

Throughput Optimal and Fast Near-Optimal Scheduling with Heterogeneously Delayed Network-State Information (Extended Version)

Srinath Narasimha, Joy Kuri and Albert Sunny

Abstract—We consider the problem of distributed scheduling in wireless networks where heterogeneously delayed information about queue lengths and channel states of all links are available at all the transmitters. In an earlier work (by Reddy et al. in *Queueing Systems*, 2012), a throughput optimal scheduling policy (which we refer to henceforth as the *R* policy) for this setting was proposed. We study the *R* policy, and examine its two drawbacks – (i) its huge computational complexity, and (ii) its non-optimal average per-packet queueing delay. We show that the *R* policy unnecessarily constrains itself to work with information that is more delayed than that afforded by the system. We propose a new policy that fully exploits the commonly available information, thereby greatly improving upon the computational complexity and the delay performance of the *R* policy. We show that our policy is throughput optimal. Our main contribution in this work is the design of two fast and near-throughput-optimal policies for this setting, whose explicit throughput and runtime performances we characterize analytically. While the *R* policy takes a few milliseconds to several tens of seconds to compute the schedule once (for varying number of links in the network), the running times of the proposed near-throughput-optimal algorithms range from a few microseconds to only a few hundred microseconds, and are thus suitable for practical implementation in networks with heterogeneously delayed information.

Index Terms—Distributed Scheduling, Heterogeneous Delays, Network State Information, Throughput Optimality

I. INTRODUCTION

A central problem in any wireless network is one of scheduling the different users in the network with the objective of optimizing some desired metric – for example, maximizing the system throughput – in the presence of challenges that are unique to the wireless medium – namely, channel fading and interference due to transmissions from other users in the network. This problem has been studied extensively in the literature. A highly influential and often cited work in this area is the work by Tassiulas and Ephremides [2], who proposed the *Back-Pressure* scheduling algorithm (a version of the *Max-Weight* algorithm [3], [4]), which is a centralized algorithm that schedules the links in the network based on global knowledge of the instantaneous queue lengths at all the links. Even though this algorithm is provably throughput optimal (see Sec. II-E), it is a centralized algorithm that requires solving a global optimization problem in each time

S. Narasimha, J. Kuri and A. Sunny are with the Department of Electronic Systems Engineering, Indian Institute of Science, Bengaluru 560012 India. e-mail: srinathn, kuri, salbert@dese.iisc.ernet.in

slot, and it also requires knowledge of instantaneous queue lengths at all links in the network to determine the schedule [1] [2].

The *Max-Weight* algorithm, being a centralized policy, involves the computationally costly task of finding a maximal independent (i.e., non-interfering) set of links that can be activated simultaneously and whose summation of link-weights is maximum [7]. To circumvent this limitation, researchers have considered two broad approaches [7] – namely, design of random-access algorithms in which access probabilities are dependent on queue sizes [7] or on arrival rates [8]–[13], and design of distributed implementations of the *Max-Weight* algorithm [14], [15]. Some of these approaches require knowledge of instantaneous queue lengths and/or instantaneous channel states to attain their objective.

Even though it may be reasonable to assume that each node has knowledge of instantaneous queue lengths and channel states (at all times) for its own links (links emanating from itself), it is less pragmatic to assume that any node possesses instantaneous information about any link in the network other than its own links (at any time instant). This could be because, for example, these quantities vary quickly with time (e.g., fast fading), or because the propagation delay of the feedback channel is large [16]. In [17], the authors consider networks where each node possesses knowledge of instantaneous queue lengths and channel states for its own links, but only has these information from other links in the network with some globally fixed delay (commonly referred to as *homogeneous delay*). The assumption of homogeneous delays, however, is not satisfied in many networks where there is often a mismatch in the delays with which each node can acquire information about queue lengths and channel states of other links in the network [1]. An example of such a case is a network where the “quality” of information that a node possesses, of queue lengths and channel states of links other than its own direct links, is a decreasing function of their “distances” from this node [1]. Non-homogeneous delays are commonly referred to as *heterogeneous delays*.

One serious issue with heterogeneous delays is that the nodes could potentially have different views of the state of the network [1]. In [1], the authors consider distributed scheduling in a wireless network where information about the queue lengths and channel states of links in the network, available at the nodes, are heterogeneously delayed. They characterize the

system throughput region for this setting, propose a scheduling policy, and show that their policy is throughput optimal. We study the limitations of the policy proposed in [1].

As in [1], we consider the problem of distributed scheduling in wireless networks where delayed queue length and delayed channel state information of each wireless link in the network, are available at all the transmitters in the network, but with possibly different delays. We refer to these as heterogeneously delayed queue state information (QSI) and heterogeneously delayed channel state information (CSI) respectively, and collectively refer to them as heterogeneously delayed network state information (NSI). We refer the reader to Sec. II-D for detailed information about the structure of heterogeneously delayed NSI that we have considered in this work. Our contributions are summarized below.

A. Our Contributions

Our primary contribution in this work is the design, analysis and performance evaluation of two fast and near-throughput-optimal scheduling policies for the case of networks with heterogeneously delayed NSI. Starting with the policy proposed in [1] (which we refer to henceforth as the R policy), we arrive at our fast and near-throughput-optimal policies through the following steps of identifying issues with the R policy and making progressive refinements to address these issues:

- 1) First, we study the limitations of the R policy. The R policy is formulated as a functional optimization problem that searches for optimal threshold functions (see Sec. III) that maximize the queue-length weighted aggregate expected throughput of the system. We study the computational complexity of the R policy and note that the R policy is computationally very costly – (i) it needs to evaluate a number of functions in its optimization domain that grows as a double exponential in the number of links in the network, and grows as an exponential in the number of levels into which the channel states on the wireless links are quantized, and (ii) it needs to consider, in computing the expected rate, a number of sample paths that is exponential in the number of links in the network and in the maximum of the heterogeneous delay values.
- 2) We show that the delay performance of the R policy is non-optimal, and that it can be improved upon significantly.
- 3) Next, we show that the structure of heterogeneously delayed NSI as defined in the system model in [1] affords each node access to NSI that is *less delayed* than that used in [1], and yet commonly available at all nodes in the network. We propose two modifications to the R policy to make use of this less delayed NSI (we call the resulting policy the H policy). We show that the H policy hugely improves upon the computational complexity of the R policy, and numerically show that the H policy also improves upon the delay performance of the R policy substantially. We establish that, like the R policy, the H policy too is throughput optimal.
- 4) Finally, we show that despite the huge leap that the H policy takes in reducing the computational complexity in comparison to that of the R policy, the computational complexity of the H policy still remains impractically large,

making the case for low-complexity scheduling policies. Taking design cues from the H policy, we propose two low-complexity near-throughput-optimal algorithms that are several orders of magnitude faster than the R and H policies. We obtain explicit analytical expressions for the expected saturated system throughputs of these policies, evaluate their performance and show that these policies closely approximate the optimal throughputs of the R and H policies. We also show that these policies possess desirable queueing delay characteristics.

II. SYSTEM MODEL

Our system model is precisely the same as that used in [1]. We use a slotted-time system and restrict our attention to single-hop transmissions. We borrow heavily in definitions, notations and nomenclature from [1] to keep our presentation easily cross-referenced and compared with [1].

A. Network Model

Our model of the wireless network has L transmitter-receiver pairs (or links); the set of links is denoted \mathcal{L} . We abstract the channel condition on each of the wireless links by the link's capacity. We model the time-varying capacity of each link l as a separate discrete-time Markov chain (DTMC) denoted $\{C_l[t]\}$ on the same state space $\mathcal{C} = \{c_1, c_2, \dots, c_M\}$, where $c_1 < c_2 < \dots < c_M$ are non-negative integers, and with the same transition probabilities^{1,2}. We assume that the channel conditions on the wireless links are all independent of each other, but identically distributed, with transition probabilities $p_{ij} := \Pr[C_l[t+1] = c_j | C_l[t] = c_i]$. We assume that these DTMCs are irreducible and aperiodic, and therefore have a stationary distribution, with the stationary probability of being in state $c_j, j \in \{1, 2, \dots, M\}$ denoted $\pi(c_j)$.

B. Interference Model

For each link l , we let I_l denote the set of links in the network that interfere with transmissions on link l ; thus, $I_l \cup \{l\}$ is an *interference set*.^{3,4} We say that a *collision* occurs with a transmission scheduled on link l if, in the same time slot, a transmission is scheduled on at least one link $l' \in I_l$. When a link l successfully transmits (i.e. when the packets transmitted on link l do not encounter a collision), a number of packets equal to $C_l[t]$ in the case of saturated queues, and equal to $\min(C_l[t], Q_l[t])$ in the case of non-saturated queues, are successfully received by the receiver on the other side of the link. However, in the case of a collision, $\gamma_l C_l[t]$ packets in the case of saturated queues, and $\min(\gamma_l C_l[t], Q_l[t])$ in the

¹We use \mathcal{C} to denote both the set and its cardinality.

²We remark that the above channel model is assumed for making notations simpler and to enhance clarity. Our results hold even for the case of networks where each link is modeled as a separate DTMC (with different state spaces and/or different transition probabilities).

³ I_l can capture arbitrary interference constraints.

⁴An *interference set* is a set of wireless links such that if there is transmission on more than one link in the set in the same time slot, then these transmissions interfere with one another, possibly resulting in only a part of the transmission being received successfully at (any of) the intended receiver(s).

TABLE I: An instance of (small) delay values for a wireless network with three links

	At TX A^\dagger	At TX B^\dagger	At TX C^\dagger
Delay in obtaining NSI of link l_1	0	1	3
Delay in obtaining NSI of link l_2	2	0	4
Delay in obtaining NSI of link l_3	1	2	0

[†] A, B, C are transmitter nodes of links l_1, l_2 , and l_3 respectively.

case of non-saturated queues, where $\gamma_l \in [0, 1]$, are received at the intended receivers.⁵

C. Traffic Model and Queue Dynamics

As noted earlier, we use a slotted-time model. Packets arriving at a transmitter, depending on their intended destination, are assigned a wireless link. Prior to transmission on this link, these packets are buffered in the queue associated with this link. We denote by $Q_l[t]$ the length of the queue associated with link l , at time t . We model the packet arrivals into the queue associated with link l , by an arrival process denoted $A_l[t]$, that describes the number of packets that arrive into the queue, at time t . For every link l , we assume that $A_l[t]$ is an integer-valued process, independent across time slots t , with $0 \leq A_l[t] \leq A_{max} < \infty$ almost surely, with $\lambda_l := \mathbb{E}[A_l[t]] < \infty$. The queue lengths are governed by the update equation:

$$Q_l[t+1] = (Q_l[t] + A_l[t] - S_l[t])^+, \quad (1)$$

where $S_l[t] := C_l[t]$ is the maximum possible service rate of link l at time t , and $x^+ := \max(0, x)$.

D. Information Model and Structure of Heterogeneously Delayed NSI

We use precisely the same structure for heterogeneously delayed NSI that each node has access to, as that used in [1]. In our model, at time t , the transmitter node of each link l has information of the current queue length and the current channel state of link l , but only has delayed queue length and delayed channel state information of other links in the network, where these delays are possibly heterogeneous but time-invariant. As an example, consider the delay values in Table I for a wireless network with three links l_1, l_2, l_3 with A, B, C as the transmitter nodes of these links respectively, such that they form a network with *complete interference*.⁶ From the first column of this table, transmitter A has NSI of link l_1 (a link emanating from transmitter A) with a delay of 0 time slots, NSI of links l_2 and l_3 with delays of 2 and 1 time slots respectively. Similarly, transmitter B has NSI of links l_1, l_2 (a link emanating from transmitter B), and l_3 with delays of 1, 0, and 2 time slots respectively, and transmitter C has NSI of links l_1, l_2 , and l_3 (a link emanating from transmitter C) with delays of 3, 4, and 0 time slots respectively. Analogously, transmitters A, B and C possess the NSI of link l_1 with delays of 0, 1, and 3 time slots respectively, of link l_2 with delays of

⁵We assume that $\{\gamma_l\}_{l \in \mathcal{L}}$ are such that $\{\gamma_l c_1, \dots, \gamma_l c_M\}_{l \in \mathcal{L}}$ for $\gamma_l \in [0, 1]$ are all integers.

⁶A network is said to have the complete interference property if each link in the network is in the interference set of all other links in the network.

2, 0, and 4 time slots respectively, and of link l_3 with delays of 1, 2, and 0 time slots respectively. Additionally, the entire table of time-invariant delay values is assumed to be available at all transmitters in the network.

We will need some notations. Let $\tau_l(h)$ denote the time-invariant delay with which the NSI of link h is available at the transmitter node of link l . For example, referring to the (2, 3) entry of Table I, $\tau_{l_3}(l_2)$ is 4. Let τ_{max} denote the maximum of the delay values across the whole network, i.e., $\tau_{max} := \max_{l, h \in \mathcal{L}, l \neq h} \tau_l(h)$. Also, let $\tau_{l,max}$ denote the maximum of the delays with which the NSI of link l are available at the transmitter node of any link in the network, i.e., $\tau_{l,max} := \max_{h \in \mathcal{L}, l \neq h} \tau_h(l)$. In other words, referring to the table of delay values, τ_{max} refers to the maximum of all the entries in the table ($\tau_{max} = 4$ in Table I), and $\tau_{l_i,max} \forall l_i \in \mathcal{L}$ refers to the maximum of the entries in row i ($\tau_{l_1,max} = 3$, $\tau_{l_2,max} = 4$, and $\tau_{l_3,max} = 2$ in Table I). We note that $\tau_l(l) = 0 \forall l \in \mathcal{L}$ in our model (the diagonal entries in Table I).

Let $C_l[t](0 : \tau) := \{C_l[t - \tau], C_l[t - \tau + 1], \dots, C_l[t]\}$, $\mathbf{C}[t](0 : \tau_{max}) := \{C_l[t](0 : \tau_{max})\}_{l \in \mathcal{L}}$ and $\mathbf{C}[t](0 : \tau_{.,max}) := \{C_l[t](0 : \tau_{l,max})\}_{l \in \mathcal{L}}$. Let $\mathbf{Q}[t](0 : \tau_{max})$ and $\mathbf{Q}[t](0 : \tau_{.,max})$ be similarly defined. We denote the NSI available at the transmitter node of link l by $\{\mathcal{P}_l(\mathbf{Q}[t](0 : \tau_{.,max})), \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{.,max}))\}$, where

$$\begin{aligned} \mathcal{P}_l(\mathbf{Q}[t](0 : \tau_{.,max})) &:= \{\mathcal{P}_{lm}(\mathbf{Q}[t](0 : \tau_{.,max}))\}_{m \in \mathcal{L}}, \\ \mathcal{P}_{lm}(\mathbf{Q}[t](0 : \tau_{.,max})) &:= \{Q_m[t - \tau]\}_{\tau=\tau_{m,max}}^{\tau_l(m)}, \\ \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{.,max})) &:= \{\mathcal{P}_{lm}(\mathbf{C}[t](0 : \tau_{.,max}))\}_{m \in \mathcal{L}}, \text{ and} \\ \mathcal{P}_{lm}(\mathbf{C}[t](0 : \tau_{.,max})) &:= \{C_m[t - \tau]\}_{\tau=\tau_{m,max}}^{\tau_l(m)} \end{aligned} \quad (2)$$

Example 1: With reference to the delay values in Table I, the NSI available at transmitter A (the transmitter node of link l_1), $\{\mathcal{P}_l(\mathbf{Q}[t](0 : \tau_{.,max})), \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{.,max}))\}$ is the set: $\{Q_{l_1}[t - 3], \dots, Q_{l_1}[t], Q_{l_2}[t - 4], \dots, Q_{l_2}[t - 2], Q_{l_3}[t - 2], Q_{l_3}[t - 1], \{C_{l_1}[t - 3], \dots, C_{l_1}[t], C_{l_2}[t - 4], \dots, C_{l_2}[t - 2], C_{l_3}[t - 2], C_{l_3}[t - 1]\}\}$.

Note that, in [1], the NSI available at link l is defined as $\{\mathcal{P}_l(\mathbf{Q}[t](0 : \tau_{max})), \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{max}))\}$, and as a consequence, even though the structure of NSI that we use in our work is the same as that in [1], the NSI at link l in our model is a subset of the corresponding NSI used in [1] since for each l , $\tau_{l,max} \leq \tau_{max}$.

It is crucial to note that $\{Q_l[t - \tau_{l,max}], C_l[t - \tau_{l,max}]\}_{l \in \mathcal{L}}$ are all common information available at the transmitter node of each link in the network, as are $\{Q_l[t - \tau_{max}], C_l[t - \tau_{max}]\}_{l \in \mathcal{L}}$. It is intuitively appealing that $\{Q_l[t - \tau_{l,max}], C_l[t - \tau_{l,max}]\}_{l \in \mathcal{L}}$, being less delayed (i.e., more recent) compared to $\{Q_l[t - \tau_{max}], C_l[t - \tau_{max}]\}_{l \in \mathcal{L}}$, and also being commonly available at all transmitters in the network, it is preferable that each transmitter in the network base its transmit/no-transmit decision on this information rather than on $\{Q_l[t - \tau_{max}], C_l[t - \tau_{max}]\}_{l \in \mathcal{L}}$. In Sec. IV we show that this intuition is indeed correct.

E. Performance Metric

The metric of interest to us is the *throughput region*⁷ (except in Sec. VII where we are interested in the saturated system throughput⁸). For this, we define the *state* of the network at time t as the process $\mathbf{Y}[t] := \{Q_l[t](0 : \tau_{l,max}), C_l[t](0 : \tau_{l,max})\}_{l \in \mathcal{L}}$. We denote $\mathbf{Y}[t]$ under the scheduling policy⁹ \mathcal{F} by $\mathbf{Y}^{\mathcal{F}}[t]$. It is easily seen that this process is a DTMC.

Given an arrival rate vector $\{\lambda_l\}_{l \in \mathcal{L}}$, we say that the network is *stochastically stable* under the scheduling policy \mathcal{F} if the DTMC $\mathbf{Y}^{\mathcal{F}}[t]$ is positive recurrent. We say that an arrival rate vector $\{\lambda_l\}_{l \in \mathcal{L}}$ is *supportable* if some scheduling policy makes the network stochastically stable for this arrival rate vector. We say that a scheduling policy is *throughput optimal* if it stabilizes the network for any arrival rate vector that is supportable.¹⁰

F. Characterization of System Throughput Region

Consider the collection of functions $\{f_l\}_{l \in \mathcal{L}}$, where $f_l : \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max})) \rightarrow \{0, 1\}$ has the following semantics – in each time slot t , the transmitter associated with each link l , computes the binary value $f_l(\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max})))$ and attempts to transmit on link l if and only if the outcome is 1. Note that the outcome of f_l is independent of all queue-length information.

Let $S_l(\mathbf{c}, \mathbf{f})$ be the expected data transmission rate (i.e. the expected number of bits or packets transmitted per time-slot) at time t that the transmitter associated with link l would receive if it applies the scheduling policy f_l (where $\mathbf{f} := \{f_l\}_{l \in \mathcal{L}}$) when the delayed CSI at time t is $\mathbf{C}[t - \tau_{l,max}] = \mathbf{c}$. That is,

$$S_l(\mathbf{c}, \mathbf{f}) = \mathbb{E} \left[C_l[t] f_l(\mathcal{P}_l(\cdot)) \left(\gamma_l + (1 - \gamma_l) \prod_{m \in I_l} (1 - f_m(\mathcal{P}_m(\cdot))) \right) \middle| \mathbf{C}[t - \tau_{l,max}] = \mathbf{c} \right], \quad (3)$$

where $\mathbf{C}[t - \tau_{l,max}] := \{C_l[t - \tau_{l,max}]\}_{l \in \mathcal{L}}$, and $\mathcal{P}_l(\cdot) := \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max}))$. Let $\mathbf{S}(\mathbf{c}, \mathbf{f}) = \{S_l(\mathbf{c}, \mathbf{f})\}_{l \in \mathcal{L}}$. Also let

$$\eta(\mathbf{c}) = \mathcal{CH}_{\mathbf{f}}(\mathbf{S}(\mathbf{c}, \mathbf{f})), \quad (4)$$

where $\mathcal{CH}_{\mathbf{f}}(\cdot)$ is the convex hull taken over all the threshold function vectors \mathbf{f} . We note that $\eta(\mathbf{c}) \subset \mathbb{R}^L$ is the convex hull of the expected data transmission rates at time t of all links in the network when $\mathbf{C}[t - \tau_{l,max}] = \mathbf{c}$, taken over all

⁷Throughput region (also called *stability region*) is the set of all *supportable* (defined in the next paragraph) mean arrival rate vectors.

⁸Saturated system throughput is the sum of the throughputs on each link in the network when the queues at the transmitter node of each link have an infinite number of packets backlogged.

⁹For each link l , a scheduling policy is a map from the NSI available at the transmitter node of link l , $\{\mathcal{P}_l(\mathbf{Q}[t](0 : \tau_{l,max})), \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max}))\}$ to a transmit/no-transmit scheduling decision for (the transmitter associated with) link l .

¹⁰It is inherent in the definition of throughput optimality that a policy that is throughput optimal need only stabilize the network for any arrival rate vector that is supportable by any policy that uses the same information structure. For example, it is possible that a policy that has access to only delayed NSI (with a given structure) to be throughput optimal and yet not support an arrival rate vector that is supportable by a policy that has access to instantaneous NSI for all links in the network.

the threshold function vectors \mathbf{f} . Let $\Lambda \subset \mathbb{R}^L$, as defined in Equation (5), be our candidate for the throughput region of the system.

$$\Lambda = \{\lambda : \lambda = \sum_{\mathbf{c} \in \mathcal{C}^L} \pi(\mathbf{c}) x(\mathbf{c}), x(\mathbf{c}) \in \eta(\mathbf{c})\} \quad (5)$$

III. BRIEF DESCRIPTION OF THE *R* POLICY

To put our work (which we develop starting in Sec. IV) in context, we briefly describe the *R* policy proposed in [1] and investigate its two drawbacks in this section. For the *R* policy, given $\mathbf{C}[t](0 : \tau_{max})$, the critical NSI of the network at time t , is defined as¹¹

$$\mathcal{CS}^R(\mathbf{C}[t](0 : \tau_{max})) := \{\{C_l[t - \tau_k(l)]\}_{k \in \mathcal{L}, k \neq l}\}_{l \in \mathcal{L}}, \quad (6)$$

and the critical NSI available at the transmitter node of link l as

$$\begin{aligned} \mathcal{CS}_l^R(\mathbf{C}[t](0 : \tau_{max})) &:= \mathcal{CS}^R(\mathbf{C}[t](0 : \tau_{max})) \\ &\cap \mathcal{P}_l^R(\mathbf{C}[t](0 : \tau_{max})), \text{ where} \end{aligned} \quad (7)$$

$\mathcal{P}_l^R(\mathbf{C}[t](0 : \tau_{max})) := \{\mathcal{P}_{lm}^R(\mathbf{C}[t](0 : \tau_{max}))\}_{m \in \mathcal{L}}$, and $\mathcal{P}_{lm}^R(\mathbf{C}[t](0 : \tau_{max})) := \{C_m[t - \tau]\}_{\tau=\tau_{max}}^{T_l(m)}$ are analogous to the corresponding definitions in Eq. (2).

In [1], the authors proposed a distributed scheduling policy for networks with heterogeneously delayed NSI. We reproduce their algorithm verbatim here. The algorithm has two steps. At each time slot,

- **Step 1:** all the transmitters compute threshold functions based on common NSI available at all transmitters. These threshold functions, one for each transmitter, map the respective transmitter's critical NSI to a corresponding threshold value, and are computed by solving the following optimization problem:

$$\arg \max_{\mathbf{T}} \sum_{l \in \mathcal{L}} Q_l[t - \tau_{max}] R_{l,\tau_{max}}(\mathbf{T}), \quad (8)$$

where

$$R_{l,\tau}(\mathbf{T}) := \mathbb{E} \left[C_l[t] \mathbf{1}_{\{C_l[t] \geq T_l(\cdot)\}} \left(\gamma_l + (1 - \gamma_l) \prod_{m \in I_l} \mathbf{1}_{\{C_m[t] < T_m(\cdot)\}} \right) \middle| \mathbf{C}[t - \tau] \right], \quad (9)$$

$$\text{and } T_l(\cdot) := T_l(\mathcal{CS}_l^R(\mathbf{C}[t](0 : \tau_{max}))).$$

- **Step 2:** each transmitter observes its current critical NSI, evaluates its threshold function (found in Step 1) at this critical NSI, and attempts to transmit if and only if its current channel rate exceeds the threshold value, i.e., when $C_l[t] \geq T_l(\mathcal{CS}_l^R(\mathbf{C}[t](0 : \tau_{max})))$.

To briefly illustrate the key idea behind the *R* policy and the difficulty that the problem of scheduling with heterogeneously delayed NSI entails, consider a network with three links –

¹¹We remark that all the quantities that pertain to the *R* policy that we show with a superscript R (e.g., $\mathcal{CS}^R(\cdot)$, $\mathcal{P}_l^R(\cdot)$) appear without superscripts in [1].

l_1 , l_2 and l_3 with perfect collision and complete interference (illustration is easier for this setting). First, let us consider the case where the transmitter node of each link has the τ -unit delayed NSI of each link (the case of homogeneous delays), and no link possesses the NSI of any link (including that of itself) with a delay lesser than τ units (this of course does not conform to the heterogeneously delayed NSI setting). In this case, the R policy can be thought of as optimally partitioning the space of all sample paths (so as to maximize the achieved system throughput) into L partitions (where some partitions are possibly empty) with the connotation that the i th partition contains sample paths for which link l_i will carry transmission when any of these sample paths is realized. To see this, consider the 3-tuple $(C_{l_1}[t-\tau], C_{l_2}[t-\tau], C_{l_3}[t-\tau])$ which give rise to eight possibilities when $\mathcal{C} = \{1, 2\}$. We can think of all the sample paths as grouped into eight classes corresponding to these eight possibilities of $(C_{l_1}[t-\tau], C_{l_2}[t-\tau], C_{l_3}[t-\tau])$. These eight classes are then partitioned optimally into three partitions such that when any sample path in the first partition is realized, link l_1 will carry transmission (and links l_2 and l_3 won't), and so on. When a particular sample path is realized, the transmitter nodes of links l_1 , l_2 and l_3 look at the tuple $(C_{l_1}[t-\tau], C_{l_2}[t-\tau], C_{l_3}[t-\tau])$ and determine the partition to which it belongs and thereby decide which of the three links will carry transmission. Now, this is true when the delays are homogeneous. In the case of heterogeneous delays, the problem of partitioning the sample paths is slightly more complicated by the fact that the different transmitters in the network have disparate views of the state of the network. We continue this illustration for the heterogeneously delayed NSI setting after introducing the example in Sec. III-A

A. Computational Complexity

From Expr. (8), we see that the domain of optimization is the set of threshold function vectors \mathbf{T} , where $\mathbf{T} = \{T_l(\cdot)\}_{l \in \mathcal{L}}$. From Expr. (9) we see that, while computing the conditional expectation, given a threshold function vector \mathbf{T} , each sample path in the evaluation of the conditional expectation is subjected to evaluation by the threshold functions $\{T_l(\cdot)\}_{l \in \mathcal{L}}$. As a consequence, there are two aspects to the computational complexity of the R policy – namely, (i) *functional evaluation complexity* – the number of threshold function vectors \mathbf{T} over which the optimization is to be done, and (ii) *sample-path complexity* – the number of sample paths that are to be considered in evaluating the conditional expectation. We first consider the complexity of functional evaluation and obtain an expression for the number of threshold function vectors required in the domain of optimization in the general case. Next, we characterize the structure on the delay values (in the table of delay values) that produces the worst-case functional evaluation complexity in the R policy. We then consider the complexity of functional evaluation in this worst-case setting, and obtain an expression for the number of threshold function vectors required in the domain of optimization in the worst case.

Proposition 1. *For the R policy, the total number of threshold functions that are needed to be considered in the domain of optimization in Expr. (8), in general, is $(\mathcal{C}+1)^{\sum_{l_i \in \mathcal{L}} \mathcal{C}^{|\mathcal{V}_i|}}$, where*

\mathcal{C} is the number of channel states, \mathcal{L} is the set of links in the network, and $\mathcal{V}_i := CS_{l_i}^R(\mathbf{C}[t](0 : \tau_{max})) \setminus \{C_{l_i}[t-\tau_{max}]\}_{l_i \in \mathcal{L}}$ is the set of input parameters to the threshold function $T_{l_i}(\cdot)$.

We present a proof of this result in Appendix A. We illustrate this result in Example 2. Next, we characterize the structure on the delay values that produces the worst-case functional evaluation complexity in the R policy.

Proposition 2. *The worst-case complexity of functional evaluation in the R policy is realized when the delays in the table of delay values are all distinct, and the delay values at positions (i, j) in row i of the table of delay values, for $j = 1, 2, \dots, i-1, i+1, \dots, L$ appear in descending order, for all i .*

We relegate the proof of this result to Appendix B. We now characterize the functional evaluation complexity of the R policy for the worst-case setting noted in Proposition 2.

Proposition 3. *For the R policy, the total number of threshold functions that are needed to be considered in the domain of optimization in Expr. (8), for the worst-case scenario noted in Proposition 2, is $(\mathcal{C}+1)^{\sum_{l_i \in \mathcal{L}} \mathcal{C}^{i(L-2)+(L-1)}}$, where \mathcal{C} is the number of channel states and L is the number of links in the network.*

From this result, the number of threshold functions required in the domain of optimization in Expr. (8) is doubly exponential in the number of links in the network, and exponential in the number of levels into which the channel states on the wireless links are quantized. We present a proof of this proposition in Appendix C. Next, we characterize the sample-path complexity of the R policy.

Proposition 4. *For the R policy, the number of sample paths that are needed to be considered in the evaluation of the conditional expectation in Expr. (9) is given by $\mathcal{C}^{L\tau_{max}}$, where \mathcal{C} is the number of channel states and L is the number of links in the network.*

From this result, the number of sample paths involved in the evaluation of the conditional expectation in Expr. (9) is exponential in the number of links in the network and in the maximum of the heterogeneous delay values. We present a proof of this proposition in Appendix D. As noted earlier, given a threshold function vector $\mathbf{T} = \{T_l(\cdot)\}_{l \in \mathcal{L}}$, each of these $\mathcal{C}^{L\tau_{max}}$ sample paths is subjected to evaluation by the L threshold functions in \mathbf{T} . We now illustrate the functional evaluation and sample-path complexities with an example.

*Example 2:*¹² Consider a three node network with three links – l_1 , l_2 , and l_3 with heterogeneous delays $\tau_{l_2}(l_1)=11, \tau_{l_3}(l_1)=7, \tau_{l_1}(l_2)=9, \tau_{l_3}(l_2)=8, \tau_{l_1}(l_3)=12, \tau_{l_2}(l_3)=6$. Note that these delay values are already in the worst-case form of Proposition 2. From these delays, we get $CS_{l_1}(\cdot)=\{C_{l_1}[t-11], C_{l_1}[t-7], C_{l_2}[t-9], C_{l_3}[t-12]\}$, $CS_{l_2}(\cdot)=\{C_{l_1}[t-11], C_{l_2}[t-9], C_{l_2}[t-8], C_{l_3}[t-12], C_{l_3}[t-6]\}$, and $CS_{l_3}(\cdot)=\{C_{l_1}[t-11], C_{l_1}[t-7], C_{l_2}[t-9], C_{l_2}[t-8], C_{l_3}[t-12], C_{l_3}[t-6]\}$. Therefore, given $\{C_l[t-\tau_{max}]\}_{l \in \mathcal{L}}$, $T_{l_1}(\cdot), T_{l_2}(\cdot), T_{l_3}(\cdot)$ are functions of 3, 4 and 5 variables respectively (i.e., in the language of Proposition 1,

¹² Reviewing the proof of Proposition 3 in Appendix C will be helpful in understanding this example fully.

$\mathcal{V}_1 = 3$, $\mathcal{V}_2 = 4$, and $\mathcal{V}_3 = 5$). Taking $\mathcal{C} = \{1, 2\}$, $T_{l_1}(\cdot)$ maps the 23 values in its domain to the real numbers say 0.5, 1.5 and 2.5, independently.¹³ Thus, the number of choices of threshold functions for $T_{l_1}(\cdot)$ is 3^8 . Similarly, the number of choices for $T_{l_2}(\cdot)$ and $T_{l_3}(\cdot)$ are 3^{16} and 3^{32} respectively. Therefore, the total number of threshold function vectors \mathbf{T} in the domain of optimization¹⁴ is 3^{56} . The number of sample paths to be considered in evaluating the conditional expectation is 2^{36} since $\tau_{max} = 12$ and $L = 3$. \square

Continuing the discussion on the key idea behind the R policy and the difficulty that the problem of scheduling with heterogeneously delayed NSI entails from the last paragraph in Sec. III, for the case of heterogeneous delays in the above example, when a particular sample path is realized, the transmitter node of link l_1 looks at $T(CS_{l_1}^R(\cdot))$ (i.e., at $T(C_{l_1}[t-11], C_{l_1}[t-7], C_{l_2}[t-9], C_{l_3}[t-12])$) and decides whether link l_1 will carry transmission depending on whether $C_{l_1}[t] \geq T(CS_{l_1}^R(\cdot))$, the transmitter node of link l_2 looks at $T(CS_{l_2}^R(\cdot))$ (i.e., at $T(C_{l_1}[t-11], C_{l_2}[t-9], C_{l_2}[t-8], C_{l_3}[t-12], C_{l_3}[t-6])$) to decide, and so on. Even though $CS_{l_1}^R(\cdot)$, $CS_{l_2}^R(\cdot)$ and $CS_{l_3}^R(\cdot)$ do not coincide, it is easy to convince oneself that the sample paths will be partitioned into three partitions as before, since we are dealing with perfect collision interference (i.e., $\gamma_l = 0$, $\forall l$).

B. The Twin DQIC Policies

With the objective of isolating the undesirable consequences of using τ_{max} -delayed queue lengths (in comparison to using $\tau_{,max}$ -delayed queue lengths) on the delay performance of the system, we consider two scheduling policies both of which have access to instantaneous CSI but differ (only) in the delays that they use to access the queue lengths. The first policy, *DQIC1* (for **D**elayed **Q**ueue lengths and **I**nstantaneous **C**hannel states), uses τ_{max} -delayed queue lengths to stabilize the queues, whereas the second policy, *DQIC2*, uses $\tau_{,max}$ -delayed queue lengths. Specifically, the two policies are:

$$\begin{aligned} \mathbf{DQIC1} &: \arg \max_l \{Q_l[t - \tau_{max}] \times C_l[t]\}_{l \in \mathcal{L}} \\ \mathbf{DQIC2} &: \arg \max_l \{Q_l[t - \tau_{l,max}] \times C_l[t]\}_{l \in \mathcal{L}} \end{aligned}$$

Remark: Note that the only difference between the two policies is that the *DQIC1* policy uses τ_{max} -delayed QSI whereas the *DQIC2* policy uses $\tau_{,max}$ -delayed QSI. Consequently, any difference in performance between the two policies is directly attributable to this difference in the delay values they use to access the QSI.

C. Delay Performance - A Preview

Towards developing a feel for what to expect in terms of the difference in delay performances of the different policies outlined in this section, we evaluate these policies through numerical simulation for the following specific instance of a two-transmitter network. We consider a network with two links l_1 and l_2 that carry transmission from their respective transmitter

¹³See footnote 12

¹⁴Incidentally, in this example, considering the delay values before rearranging as noted in Proposition 2, the total number of threshold function vectors \mathbf{T} in the domain of optimization is also 3^{56} .

TABLE II: An instance heterogenous delay values for a wireless network with two links

	At TX A^\dagger	At TX B^\dagger
Delay in obtaining NSI of link l_1	0	1
Delay in obtaining NSI of link l_2	x	0

[†] A, B are transmitter nodes of links l_1 and l_2 respectively.

nodes A and B to their respective destinations (with $I_{l_1} = \{l_2\}$ and $I_{l_2} = \{l_1\}$), with the channels on these links modeled as independent DTMCs on the state space $\mathcal{C} = \{1, 2\}$ with crossover probability 0.1. We consider heterogeneous delays as noted in Table II, where $x \geq 1$ is a parameter we vary. From this table, we see that $\tau_{max} = x$, $\tau_{l_1,max} = 1$, $\tau_{l_2,max} = x$.

The *DQIC1* and *DQIC2* policies specialize to this setting as follows:

$$\mathbf{DQIC1} : \arg \max \{Q_{l_1}[t - x] \times C_{l_1}[t], Q_{l_2}[t - x] \times C_{l_2}[t]\}$$

$$\mathbf{DQIC2} : \arg \max \{Q_{l_1}[t - 1] \times C_{l_1}[t], Q_{l_2}[t - x] \times C_{l_2}[t]\}$$

i.e., *DQIC1* schedules link l_1 if $Q_{l_1}[t - x] \times C_{l_1}[t] \geq Q_{l_2}[t - x] \times C_{l_2}[t]$, and link l_2 otherwise, whereas *DQIC2* schedules link l_1 if $Q_{l_1}[t - 1] \times C_{l_1}[t] \geq Q_{l_2}[t - x] \times C_{l_2}[t]$, and link l_2 otherwise.

Packets arrive into the two queues at the two transmitters as independent Poisson processes with rates $\lambda_1 = \lambda_2 = \frac{1}{4}$ (for this setting, it is easily seen that sum rates (i.e. $\lambda_1 + \lambda_2$) of up to 1.75 are supportable.¹⁵) Packets are time-stamped on arrival and on exit, and on servicing link l successfully, a number of packets equal to $\min(C_l[t], Q_l[t])$ are removed from the queue of link l . Fig. 1(a) depicts the average queueing delay per-packet (in units of time-slots) of the *DQIC1*, *DQIC2*, *R* and *H* policies. Fig. 1(b) depicts the percentage reduction in the average queueing delay per-packet in the *DQIC2* policy compared to that of the *DQIC1* policy.

Conclusion: The average per-packet queueing delays of the *DQIC1* and *R* policies grow linearly with τ_{max} whereas that of the *DQIC2* policy tends to become almost impervious to increase in τ_{max} . Comparing the *DQIC1*, *DQIC2* and *R* policies and their delay performances, the fact that the *DQIC1* and *R* policies use τ_{max} -delayed queue lengths isolates itself as the sole reason for their undesirable delay performances.

IV. THE *H* POLICY

In Sec. II-D, we conjectured that $\{C_l[t - \tau_{l,max}]\}_{l \in \mathcal{L}}$, being less delayed compared to $\{C_l[t - \tau_{max}]\}_{l \in \mathcal{L}}$, and being commonly available at all transmitters in the network, it may be preferable that each transmitter base its transmit/no-transmit decision on this information rather than on $\{C_l[t - \tau_{max}]\}_{l \in \mathcal{L}}$. In addition to testing this hypothesis, we wish to reduce the computational complexity and the average per-packet queueing delay in comparison with that of the *R* policy. Motivated by

¹⁵ Considering the four possibilities of data rates for $C_{l_1}[t]$ and $C_{l_2}[t]$, a data rate of 1 unit is supportable when $C_{l_1}[t] = C_{l_2}[t] = 1$ which happens with steady-state probability 0.25, and a data rate of 2 units are supportable in all the other three cases since when the queues are saturated, both the *DQIC1* and *DQIC2* policies would choose the link with the largest data rate. These three cases happen with steady-state probability 0.25 each.

¹⁶The *LC-ELDR* and *LC-ERDMC* policies are identical for a network with only two links (see Algorithm 1 in Sec. VI).

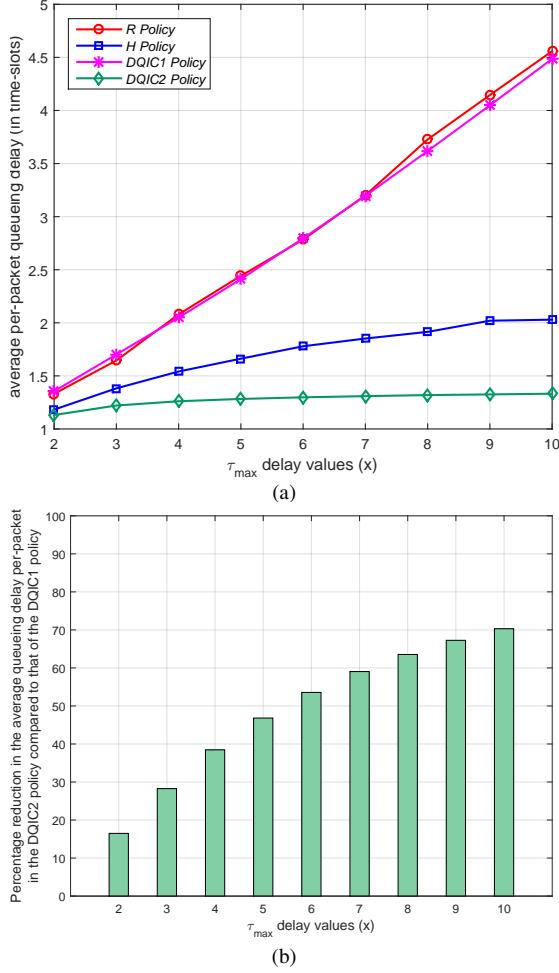


Fig. 1: (a) Comparison of average queueing delay per-packet in the DQIC1, DQIC2, R, and H policies. The R and DQIC1 policies use τ_{max} -delayed QSI whereas the H and DQIC2 policies use $\tau_{l,max}$ -delayed QSI. The average per-packet queueing delays of the DQIC1 and R policies grow linearly with τ_{max} whereas that of the DQIC2 policy tends to become almost unaffected by increase in τ_{max} . Comparing the DQIC1, DQIC2 and R policies and their delay performances, the fact that the DQIC1 and R policies use τ_{max} -delayed queue lengths isolates itself as the sole reason for their undesirable delay performances. The (relatively) small difference in the queueing delay performances of the H and DQIC2 policies is due to the fact that the DQIC2 policy has access to instantaneous CSI whereas the H policy only has access to delayed CSI. (b) Percentage reduction in the average queueing delay per-packet in the DQIC2 policy compared to that in the DQIC1 policy.

these, we define the H policy to be the following scheduling policy. In each time slot,

- **Step 1:** using the common NSI available at all transmitters, each transmitter determines the optimal threshold function vector, by solving the following optimization problem:

$$\arg \max_{\mathbf{T}} \sum_{l \in \mathcal{L}} Q_l[t - \tau_{l,max}] R_{l,\tau_{l,max}}(\mathbf{T}), \quad (10)$$

where

$$R_{l,\tau}(\mathbf{T}) := \mathbb{E} \left[C_l[t] \mathbf{1}_{\{C_l[t] \geq T_l(\cdot)\}} \left(\gamma_l + (1 - \gamma_l) \prod_{m \in I_l} \mathbf{1}_{\{C_m[t] < T_m(\cdot)\}} \right) \middle| \mathbf{C}[t - \tau] \right], \quad (11)$$

$$\begin{aligned} \mathbf{C}[t - \tau_{l,max}] &:= \{C_l[t - \tau_{l,max}]\}_{l \in \mathcal{L}}, \\ \mathcal{CS}_l(\mathbf{C}[t](0 : \tau_{l,max})) &:= \mathcal{CS}(\mathbf{C}[t](0 : \tau_{l,max})) \\ &\cap \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max})), \end{aligned} \quad (12)$$

$$\text{and } \mathcal{CS}(\mathbf{C}[t](0 : \tau_{l,max})) := \{\{C_l[t - \tau_k(l)]\}_{k \in \mathcal{L}, k \neq l}\}_{l \in \mathcal{L}}. \quad (13)$$

- **Step 2:** each transmitter observes its current critical NSI, evaluates its threshold function (found in Step 1) at this critical NSI, and attempts to transmit if and only if its current channel rate exceeds the threshold value, i.e., when $C_l[t] \geq T_l(\mathcal{CS}_l(\mathbf{C}[t](0 : \tau_{l,max})))$.

Example 3: Continuing Example 1, for the H policy, the critical NSI of the network $\mathcal{CS}(\mathbf{C}[t](0 : \tau_{l,max}))$ is the set $\{C_{l_1}[t-1], C_{l_1}[t-3], C_{l_2}[t-2], C_{l_2}[t-4], C_{l_3}[t-1], C_{l_3}[t-2]\}$ and the critical NSI available at the transmitter nodes of links l_1, l_2, l_3 , namely $\mathcal{CS}_{l_1}(\mathbf{C}[t](0 : \tau_{l,max})), \mathcal{CS}_{l_2}(\mathbf{C}[t](0 : \tau_{l,max})), \mathcal{CS}_{l_3}(\mathbf{C}[t](0 : \tau_{l,max}))$, respectively are $\{C_{l_1}[t-1], C_{l_1}[t-3], C_{l_2}[t-2], C_{l_2}[t-4], C_{l_3}[t-1], C_{l_3}[t-2]\}, \{C_{l_1}[t-1], C_{l_1}[t-3], C_{l_2}[t-2], C_{l_2}[t-4], C_{l_3}[t-2]\}$, and $\{C_{l_1}[t-3], C_{l_2}[t-4], C_{l_3}[t-1], C_{l_3}[t-2]\}$. \square

Remark: It is to be noted that the R and H policies differ only in the delay values that they use to access queue lengths and channel state information, with the R policy using the delay value τ_{max} and the H policy using the delay value $\tau_{l,max}$ to access the NSI of link l .

A. Computational Complexity

In this section, similar to the way that we dealt with the R policy in Sec. III-A, we first consider the complexity of functional evaluation and obtain an expression for the number of threshold function vectors required in the domain of optimization in the general case. Next, we characterize the structure on the delay values (in the table of delay values) that produces the worst-case functional evaluation complexity in the H policy. We then consider the complexity of functional evaluation in this worst-case setting, and obtain an expression for the number of threshold function vectors required in the domain of optimization in the worst case.

Proposition 5. For the H policy, the total number of threshold functions that are needed to be considered in the domain of optimization in Expr. (10), in general, is $(\mathcal{C} + 1)^{\sum_{l_i \in \mathcal{L}} \mathcal{C}^{|\mathcal{W}_i|}}$, where \mathcal{C} is the number of channel states, \mathcal{L} is the set of links in the network, and $\mathcal{W}_i := \mathcal{CS}_{l_i}(\mathbf{C}[t](0 : \tau_{l,max})) \setminus \{C_{l_i}[t - \tau_{l_i,max}]\}_{l_i \in \mathcal{L}}$ is the set of input parameters to the threshold function $T_{l_i}(\cdot)$.

We present a proof of this result in Appendix E. We illustrate this result in Example 4. Next, we present the following result that will be required later:

Proposition 6. $\mathcal{CS}_l^R(\mathbf{C}[t](0 : \tau_{max})) = \mathcal{CS}_l(\mathbf{C}[t](0 : \tau_{max}))$.

We provide a proof of this proposition in Appendix F. We next characterize the structure on the delay values that brings

out the worst-case functional evaluation complexity in the H policy.

Proposition 7. *The worst-case complexity of functional evaluation in the H policy is realized when the delays in the table of delay values are all distinct, and the delay values at positions (i, j) in row i of the table of delay values, for $j = 1, 2, \dots, i-1, i+1, \dots, L$ appear in descending order, for all i .*

We highlight a minor difference in the proof of this proposition compared to that of Proposition 2, in Appendix G. We now characterize the functional evaluation complexity of the H policy for the worst-case setting noted in Proposition 7.

Proposition 8. *For the H policy, the total number of threshold functions that are needed to be considered in the domain of optimization in Expr. (10), for the worst-case scenario noted in Proposition 7, is $(C+1)^{\sum_{l \in \mathcal{L}} C^{i(L-2)}}$, where C is the number of channel states and L is the number of links in the network.*

We present a proof of this result in Appendix H. Next, we characterize the sample-path complexity of the H policy.

Proposition 9. *For the H policy, the number of sample paths that are needed to be considered in the evaluation of the conditional expectation in Expr. (11) is given by $C^{\sum_{l \in \mathcal{L}} \tau_{l,\max}}$, where C is the number of channel states.*

The proof of this proposition is similar to the proof of Proposition 4 and is therefore omitted. We now illustrate with an example the reduction in computational complexity that the H policy achieves over that of the R policy.

Example 4: Consider the same network as in Example 2. After computing $CS_{l_1}(\cdot)$, $CS_{l_2}(\cdot)$ and $CS_{l_3}(\cdot)$ using Eq. (12), we see that given $\{C_l[t - \tau_{l,\max}]\}_{l \in \mathcal{L}}$, $T_{l_1}(\cdot)$, $T_{l_2}(\cdot)$, $T_{l_3}(\cdot)$ are functions of 1, 2 and 3 variables respectively. Using the same arguments as in the proofs of Propositions 3 and 8 in Appendices C and H, the number of choices of threshold functions for $T_{l_1}(\cdot)$, $T_{l_2}(\cdot)$ and $T_{l_3}(\cdot)$ are 3^2 , 3^4 and 3^8 respectively. Therefore, the total number of threshold function vectors \mathbf{T} in the domain of optimization is 3^{14} (the R policy required 3^{56} threshold functions; see Example 2) and the number of sample paths required to be considered in evaluating the conditional expectation is 2^{32} (the R policy required considering 2^{36} sample paths; see Example 2), yielding an enormously massive reduction in computational complexity over that of the R policy. We note that this vast reduction is mainly due to the fact that for the H policy, the exponent in the double exponential in the expression in Proposition 8 is smaller than that of the R policy.

B. Delay Performance

Fig. 1 shows the average per-packet queueing delay (in units of time-slots) in the H policy for the setting of the example in Sec. III-C. The average per-packet queueing delay of the R policy grows linearly with τ_{\max} whereas that of the H policy tends to flatten out. Comparing the R , H , $DQIC1$ and $DQIC2$ policies and their delay performances, it is clearly evident that the use of $\tau_{l,\max}$ -delayed queue lengths in the H and $DQIC2$ policies is what gives these policies their better delay performances in comparison to those of the R and $DQIC1$ policies which use τ_{\max} -delayed queue lengths.

C. Throughput Optimality

In this section we show that, like the R policy, the H policy too is throughput optimal. First, we prove that if an arrival process is supportable, then the expected arrival rate of this process should lie within the system throughput region Λ defined in Sec. II-F. This would then mean that the system throughput region Λ is the region that encompasses all supportable arrival rates given the NSI structure in Sec. II-D. Next, we show that the H policy stabilizes all arrival rate vectors in the system throughput region Λ . These two together would then imply that the H policy is throughput optimal.

Lemma 1. *Under the NSI structure noted in Sec. II-D, if the traffic arrival process $\{A[t]\}_t$ is supportable, then $E[A[t]] \in \Lambda$.*

The proof of this lemma, which is similar to the proof of Lemma 4.1 in [1], is available in Appendix I.

Corollary 1. *The system throughput regions defined in Equation (5), and in Sec. 4.1 in [1], are identical.*

This is an immediate consequence of Lemma 4.1 in [1] and Lemma 1 stated above. See Sec. 4.1 in [1] for the definition of the system throughput region considered in [1].

Theorem 1. *The H policy is throughput optimal.*

The proof of this theorem, which is similar to the proof of the corresponding theorem for the R policy in [1] (but with significantly more details) is available in Appendix J. *An important implication of this theorem is that using $\tau_{l,\max}$ -delayed NSI instead of τ_{\max} -delayed NSI (for each link l) does not harm the achievable system throughput, a fact that we exploit crucially in designing our computationally efficient near-throughput-optimal policies in Sec. VI.*

V. ANALYTICAL CHARACTERIZATION OF DELAY PERFORMANCE

In this section, we characterize the average queueing delay per packet in the $DQIC1$, $DQIC2$, and R policies analytically. To do this, we first fix a time-horizon, consider all possible packet arrival patterns and all possible channel states in each slot within this time horizon and compute the average delay experienced by a packet in each of the above scheduling policies in the limit as the time-horizon tends to infinity. We need the following additional notation:

$T :=$ Time horizon (i.e, the number of time slots over which we measure the average delay per packet)

$N_i :=$ Number of packets that have been serviced from time $t = 0$ to $t = T$ from the queue associated with link l_i

$N := \sum_{i=1}^{\mathcal{L}} N_i$ is the total number of packets in the system that have been serviced from time $t = 0$ to $t = T$ from the queue associated with any link in the network

$$\begin{aligned}
a_{l,j} &:= \begin{cases} \text{The arrival time of the } j\text{th packet from the head of line in the queue associated with link } l & \text{if } j \leq Q_l[t] \\ \infty & \text{otherwise} \end{cases} \\
\mathbf{1}_{\{a\}} &:= \begin{cases} 1 & \text{if } a \text{ is true} \\ 0 & \text{otherwise} \end{cases} \\
\mathbf{1}_{\{a,b\}} &:= \mathbf{1}_{\{a\}} \times \mathbf{1}_{\{b\}} \\
x^+ &:= \max\{0, x\}
\end{aligned}$$

With these definitions in place, and with the assumption that the queues are purged of the packets that were transmitted in a particular time slot at the end of that time slot, we are now ready to characterize the average per-packet queueing delay in these policies. First, we characterize $\bar{D}_{\text{arrivals,CSI}}$ – the long-term average queueing delay conditioned on the knowledge of the channel states of all links in the network for all time slots (i.e., given $\{C_l[t]\}_{l \in \mathcal{L}}$, for all t such that $-\tau_{\max} \leq t \leq T-1$) and conditioned on the availability of the actual arrivals at all links in the network for all time slots (i.e., $\{A_l[t]\}_{l \in \mathcal{L}}$, for all t such that $-\tau_{\max} \leq t \leq T-1$) as follows:

$$\bar{D}_{\text{arrivals,CSI}} := \lim_{T \rightarrow \infty} \sum_{t=-\tau_{\max}}^{T-1} \sum_{i=1}^{\mathcal{L}} \mathbf{1}_{\{\cdot\}} \sum_{j=1}^{C_{l_i}[t]} (t - a_{l_i,j})^+ \times \frac{1}{N}, \quad (14)$$

where the indicator function in Eq. (14) indicates the condition under which link l_i is scheduled for transmission in time slot t , and is defined, for the various policies, as follows:¹⁶

For the *DQIC1* Policy:

$$\mathbf{1}_{\{\cdot\}} := \begin{cases} \mathbf{1}_{\{Q_{l_i}[t] \times C_{l_i}[t] > \max_{k < i} \{Q_{l_k}[t] \times C_{l_k}[t]\}, \\ Q_{l_i}[t] \times C_{l_i}[t] \geq \max_{k > i} \{Q_{l_k}[t] \times C_{l_k}[t]\}\}} & \text{if } t < 0 \\ \mathbf{1}_{\{Q_{l_i}[t-\tau_{\max}] \times C_{l_i}[t] > \max_{k < i} \{Q_{l_k}[t-\tau_{\max}] \times C_{l_k}[t]\}, \\ Q_{l_i}[t-\tau_{\max}] \times C_{l_i}[t] \geq \max_{k > i} \{Q_{l_k}[t-\tau_{\max}] \times C_{l_k}[t]\}\}} & \text{if } t \geq 0 \end{cases}$$

In the indicator function in the expressions above, we note that the splitting of the comparison of the product of queue-length and channel state on link l_i with that on the other links, into two – namely, (i) comparison with that on links $l_k, k < i$, and (ii) comparison with that on links $l_k, k > i$, creates a lexicographic ordering among the links. This is required to consistently resolve the “winner” in case there is a tie in the queue-length channel-state product on multiple links – we always resolve in favor of the smallest i (as a convention) in case of a tie.

¹⁶ Note that the time index t starts from $-\tau_{\max}$ in Eq. (14) since at time $t = 0$, τ_{\max} -delayed NSI would be the NSI at time $t = -\tau_{\max}$. Note that N_i (and hence N) does not include the packets serviced in time slots $t = -\tau_{\max}$ to $t = -1$.

For the *DQIC2* Policy:

$$\mathbf{1}_{\{\cdot\}} := \begin{cases} \mathbf{1}_{\{Q_{l_i}[t] \times C_{l_i}[t] > \max_{k < i} \{Q_{l_k}[t] \times C_{l_k}[t]\}, \\ Q_{l_i}[t] \times C_{l_i}[t] \geq \max_{k > i} \{Q_{l_k}[t] \times C_{l_k}[t]\}\}} & \text{if } t < 0 \\ \mathbf{1}_{\{Q_{l_i}[t-\tau_{l,\max}] \times C_{l_i}[t] > \max_{k < i} \{Q_{l_k}[t-\tau_{l,\max}] \times C_{l_k}[t]\}, \\ Q_{l_i}[t-\tau_{l,\max}] \times C_{l_i}[t] \geq \max_{k > i} \{Q_{l_k}[t-\tau_{l,\max}] \times C_{l_k}[t]\}\}} & \text{if } t \geq 0 \end{cases}$$

For the *R* Policy: $\mathbf{1}_{\{\cdot\}} := \mathbf{1}_{\{C_{l_i}[t] \geq T_{l_i} (CS_{l_i}^R(\mathbf{C}[t](0:\tau_{\max})))\}}$

For the *H* Policy: $\mathbf{1}_{\{\cdot\}} := \mathbf{1}_{\{C_{l_i}[t] \geq T_{l_i} (CS_{l_i}(\mathbf{C}[t](0:\tau_{\max})))\}}$

Note that the knowledge of the number of packet arrivals into the queue associated with link l_i in each time slot t , for $-\tau_{\max} \leq t \leq T-1$ is subsumed in the $a_{l_i,j}$ term and the knowledge of the channel state on link l_i in each time slot t , for $-\tau_{\max} \leq t \leq T-1$ is accounted for in the fact that the queues are purged of the packets transmitted in a time slot at the end of that time slot, which in turn affects the $a_{l_i,j}$ term. Also note that any packets remaining in the queues after time slot $T-1$ are not serviced and hence do not contribute to the average queueing delay.

Now, removing the conditioning on knowledge of CSI, we get the expression for $\bar{D}_{\text{arrivals}}$ – the average per-packet queueing delay (conditioned only on the arrivals) as follows:

$$\bar{D}_{\text{arrivals}} := \sum_{\{c_l[t]\}_{l \in \mathcal{L}}} \bar{D}_{\text{arrivals,CSI}} \prod_{t=-\tau_{\max}}^{T-1} \left(\prod_{m=1}^{\mathcal{L}} \left(\pi(c_m[-\tau_{\max}]) \prod_{n=-\tau_{\max}}^{T-2} p_{c_m[n] c_m[n+1]} \right) \right) \quad (15)$$

where $\pi(c)$ is the steady-state probability of being in state c in the channel-state Markov chain, and p_{ij} is the one-step transition probability from state i to state j in the channel-state Markov chain.

Finally, removing the conditioning on the knowledge of arrivals on each link in each time slot t , $-\tau_{\max} \leq t \leq T-1$, we get the expression for \bar{D} – the required average per-packet queueing delay as follows:

$$\bar{D} := \sum_{\{a_l[t]\}_{l \in \mathcal{L}}} \bar{D}_{\text{arrivals}} \prod_{t=-\tau_{\max}}^{T-1} \left(\prod_{l=1}^{\mathcal{L}} \Pr(A_l[t] = a_l[t]) \right) \quad (16)$$

A. Computational Complexity

Since Expr. (16) is not in closed form, we will have to resort to numerical evaluation of this expression for a fixed finite value of T . For a fixed finite T , the number of time slots over which Expr. (16) needs to be evaluated is $\tau_{\max} + T$. At each link in the network, the number of packet arrivals in each time slot can be any of the values in the set $\{0, 1, \dots, A_{\max}\}$, where A_{\max} is the maximum number of packets that can arrive in each time slot (see Sec. II-C). Thus, the number

of possible arrival streams at a link (over all the $\tau_{max} + T$ time slots), is given by $(A_{max} + 1)^{(\tau_{max}+T)}$, and therefore the number of possible arrival streams taking all the links together is $(A_{max} + 1)^{L(\tau_{max}+T)}$. Further, in case of Markov chains where any state can be reached from any other state in one step (as in the setting considered in Sec. VII-C), the channel state on each link can be any of the \mathcal{C} states, where \mathcal{C} is the number of states in the channel Markov chain. Thus, for each possible arrival stream pattern at each link, over the fixed finite time horizon T , the number of possible channel state transition patterns at a link is given by $\mathcal{C}^{(\tau_{max}+T)}$, and therefore the number of possible channel state transition patterns taking all the links together is $\mathcal{C}^{L(\tau_{max}+T)}$. Thus, using Eq. (16), the computational complexity of comprehensively evaluating the average queueing delay per-packet over a fixed finite T is $(A_{max} + 1)^{L(\tau_{max}+T)} \times \mathcal{C}^{L(\tau_{max}+T)}$ or $((A_{max} + 1)\mathcal{C})^{L(\tau_{max}+T)}$ operations. Thus, for example, in a network with 2 transmitters, where in each time slot, the number of packet arrivals can only be 0 or 1, and where the channel state Markov chain has only 2 states, for a time horizon of $T = 10$, and for $\tau_{max} = 2$, the computational complexity is 2^{48} (or, roughly 10^{14}) operations.

VI. LOW COMPLEXITY SCHEDULING POLICIES

The H policy, despite the immense reduction in computational complexity that it achieves in comparison to the R policy, is still computationally complex, and therefore impractical. In this section, we propose and evaluate two fast and near-throughput-optimal scheduling policies – namely, *LC-ELDR* (for Eliminate link with Least Data Rate), and *LC-ERDMC* (for Eliminate link that Reduces Delays for Maximum number of Channels) – for the heterogeneously delayed NSI setting, confining our attention to the more pragmatic case of perfect collision interference (i.e., $\gamma_l = 0$, $\forall l \in \mathcal{L}$). Initially, we consider the case of complete interference (i.e., $I_l = \mathcal{L} \setminus \{l\}$ $\forall l \in \mathcal{L}$), and consider extension to the case of multiple interference sets subsequently. These low-complexity scheduling policies derive their main idea from the H policy; the idea being that the common delay value that all the contending transmitters (i.e., transmitters corresponding to contending links) can use to access the NSI of some link l (of at least one link), could be reduced (and hence more recent [i.e., more reliable] delayed link-statistics could be made use of in computing the schedule) by carefully choosing to eliminate one link from among the contending links. These policies are iterative (we call each iteration, a “round”) and they operate by choosing to eliminate one link in each round; the two policies differ only in the criterion they use for selecting the link that they eliminate in each round. Further, eliminating one link in each round accords these policies polynomial running times as we show in Sec. VI-C. We list the *LC-ELDR* policy in Algorithm 1 and demonstrate its working in detail in Sec. VI-A. The *LC-ERDMC* policy is identical to the *LC-ELDR* policy listed in Algorithm 1 except for step 16 which is listed separately in Algorithm 2. We derive analytical expressions for the exact expected saturated system throughputs of the *LC-ELDR* and *LC-ERDMC* policies in Sec. VI-B, and plot

TABLE III: Heterogeneous delays for the illustration in Sec. VI-A

	At TX A	At TX B	At TX C	At TX D
Delay in obtaining NSI of link l_1	0	4	1	1
Delay in obtaining NSI of link l_2	1	0	1	2
Delay in obtaining NSI of link l_3	1	1	0	5
Delay in obtaining NSI of link l_4	3	1	1	0

these expressions in Sec. VII and show that these policies are near-optimal.

A. Dynamics of the *LC-ELDR* Policy

We now demonstrate the working of the *LC-ELDR* policy for a network with four wireless links – l_1, l_2, l_3, l_4 (with A, B, C, D as the transmitter nodes on these links, respectively), all contending for transmission in the current time slot. We will assume that the queues at the transmitter node of these links are saturated, and that the heterogeneous delays are as in Table III. Let each of the wireless links be modeled as a Markov chain on the state space $\mathcal{C} = \{1, 2\}$ with crossover probability 0.1 (the case of *very slow varying channel*). The n -step transition probability matrices for this Markov chain are as below (shown for only those delays [n-values] that are required in our computation):

$$\begin{aligned} P^{(1)} &= \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix}, \quad P^{(2)} = \begin{bmatrix} 0.82 & 0.18 \\ 0.18 & 0.82 \end{bmatrix}, \\ P^{(3)} &= \begin{bmatrix} 0.756 & 0.244 \\ 0.244 & 0.756 \end{bmatrix}, P^{(4)} = \begin{bmatrix} 0.7048 & 0.2952 \\ 0.2952 & 0.7048 \end{bmatrix}, \\ P^{(5)} &= \begin{bmatrix} 0.6638 & 0.3362 \\ 0.3362 & 0.6638 \end{bmatrix} \end{aligned}$$

We assume the following delayed channel-realizations for our illustration: $C_{l_1}[t-4] = 2, C_{l_1}[t-1] = 1, C_{l_2}[t-2] = 2, C_{l_2}[t-1] = 1, C_{l_3}[t-5] = 2, C_{l_3}[t-1] = 2, C_{l_4}[t-3] = 2, C_{l_4}[t-1] = 1$. We first illustrate the working of the *LC-ELDR* policy when it is executed at transmitter C , and towards the end we mention the minor differences in the dynamics when it is executed at the other transmitters. Since we are initially concerned with the case of complete interference, we set $ActiveSet = \{l_1, l_2, l_3, l_4\}$ before invoking Algorithm 1, where $ActiveSet$ is the set of links that are (still) contending for transmission in the current time slot.

(step 2) transmitter C sets $I = \{l_3\}$

ROUND 1:¹⁸

(step 4) transmitter C computes $\tau_{l_i,max}$, for $i = 1, 2, 3, 4$, which results in $\tau_{l_1,max} = 4, \tau_{l_2,max} = 2, \tau_{l_3,max} = 5, \tau_{l_4,max} = 3$. This is illustrated in the last column of the table in Fig. 2(a)

(step 5) transmitter C computes the expected data rate on each link l_i in $ActiveSet$ given the $\tau_{l_i,max}$ -delayed channel-state realization of that link, as follows:¹⁹ $E[C_{l_1}[t] | C_{l_1}[t -$

¹⁷the node where this algorithm is being run

¹⁸We will call the “while” loop body from line 3 to line 23, and also the computation in lines 24 – 29 in Algorithm 1, a “round”.

¹⁹we ignore the multiplication of the delayed queue-lengths with the expected data rates since we are considering saturated queues for this illustration.

Figure 2 consists of eight sub-tables (a-h) illustrating the dynamics of the LC-ELDR policy. Each sub-table has columns for 'At TX A', 'At TX B', 'At TX C', 'At TX D', and 'τ_{i,max}'. The rows represent four links: l₁, l₂, l₃, and l₄. The values in the table represent delay values and maximum delays. The tables show the progression of the algorithm, with certain cells being highlighted in green, red, or hatched to indicate specific states or changes in the elimination process.

Fig. 2: Illustration of the dynamics of the *LC-ELDR* policy. In the beginning, $ActiveSet = \{l_1, l_2, l_3, l_4\}$. Subfigure (a) shows the original table of delay values and $\tau_{i,max}$ values computed in step 4 of Algorithm 1 in round 1. Subfigures (b), (c), (d) show the computation of the set EC in round 1 (step 7). In subfigure (b), after setting aside the link with the largest expected data rate (link l_2 ; shown by highlighting in green the column of transmitter B [the transmitter node of link l_2]), we see whether eliminating link l_1 reduces the $\tau_{i,max}$ delays for any other link in $ActiveSet$. This is done by temporarily masking the row and column pertaining to link l_1 (shown hatched), recomputing the $\tau_{i,max}$ delays and comparing with previously computed values. The value of $\tau_{l_4,max}$ reduces (from 3 to 1), and hence link l_1 is a candidate for elimination, and therefore belongs to EC . Similarly, in subfigure (c), we determine that link l_3 does not belong to EC , and in subfigure (d), we determine that link l_4 belongs to EC . In subfigure (e), we eliminate link l_1 (shown by striking-through in red the row and column pertaining to link l_1) since only l_1 and l_4 belong to EC and l_1 has smaller expected data rate than l_4 (steps 16, 21). This reduces $ActiveSet$ to $\{l_2, l_3, l_4\}$. Subfigures (f), (g), (h) show similar calculations for round 2, at the end of which l_4 is eliminated, leaving behind only l_2 and l_3 in $ActiveSet$.

$$4] = 2] = 1.7048, E[C_{l_2}[t] | C_{l_2}[t-2] = 2] = 1.82, E[C_{l_3}[t] | C_{l_3}[t-5] = 2] = 1.6638, E[C_{l_4}[t] | C_{l_4}[t-3] = 2] = 1.756$$

(step 6) transmitter C sets aside link l_2 by setting $H = l_2$, since link l_2 has the largest expected data rate in this round, as computed in the previous step (this is done so that the link with the largest expected data rate in this round is not eliminated). This is shown by highlighting transmitter B 's (B is the transmitter node on link l_2) column in green in Fig. 2(b).

(step 7) transmitter C , to decide on the link it wants to eliminate in this round, computes EC as follows:

- first, transmitter C temporarily masks the row and column pertaining to link l_1 (shown by hatching the corresponding row and column, in Fig. 2(b)), computes the $\tau_{l_i,max}$ values for links l_i , $i = 2, 3, 4$ (i.e., for all links other than link l_1), compares them to the previously computed $\tau_{l_i,max}$ values for these links, finds that upon masking the row and column pertaining to link l_1 , the $\tau_{l_i,max}$ value for at least one link (namely, that of link l_4) reduces (from 3 to 1; see Fig. 2(b)), and hence decides that link l_1 is a candidate for elimination, and adds it to the set of elimination candidates, EC . What this means is that if transmitter C does choose to eliminate link l_1 , each of the transmitters corresponding to the remaining links would then be able to access the delayed channel state of link l_4 at a reduced common delay value of 1 unit (instead of 3 units, as was required previously) since the other transmitters assuredly possess the delayed channel-state of link l_4 at this new reduced delay of 1 unit
- transmitter C skips this procedure for link l_2 since l_2 , being the link with the largest expected data rate, was set aside in the previous step
- next, transmitter C temporarily masks the row and column pertaining to link l_3 (i.e., its own link; shown again by hatching the row and column of link l_3 , in Fig. 2(c)), computes the $\tau_{l_i,max}$ values for links l_i for $i = 1, 2, 4$ (i.e., for all links other than link l_3), compares them to

the previously computed $\tau_{l_i,max}$ values for these links, finds that upon masking the row and column pertaining to link l_3 , the $\tau_{l_i,max}$ value for none of the links l_i for $i = 1, 2, 4$ reduce, and hence decides that link l_3 (i.e., its own link) is not a candidate for elimination

- lastly, transmitter C temporarily masks the row and column pertaining to link l_4 (shown by hatching the row and column of link l_4 , in Fig. 2(d)), computes the $\tau_{l_i,max}$ values for links l_i for $i = 1, 2, 3$ (i.e., for all links other than link l_4), compares them to the previously computed $\tau_{l_i,max}$ values for these links, finds that upon masking the row and column pertaining to link l_4 , the $\tau_{l_i,max}$ value for at least one link (namely, that of links l_2 and l_3) reduces (from 2 to 1 for link l_2 , and from 5 to 1 for link l_3 ; see Fig. 2(d)), and hence decides that link l_4 is a candidate for elimination, and adds it to the set of elimination candidates, EC . This means that if transmitter C does choose to eliminate link l_4 , it is assured that each of the links remaining in $ActiveSet$ would then be able to access the channel state of links l_2 and l_3 at reduced delays of 1 unit and 1 unit respectively (instead of 2 and 5 units respectively, as was required previously)

Thus, transmitter C has computed the set of elimination candidates to be $EC = \{l_1, l_4\}$

(step 16) Among the elimination candidates in set EC , transmitter C chooses link l_1 to be the one that will be eliminated in this round by setting $S = l_1$, since link l_1 has the lowest expected data rate among l_1 and l_4 (which have expected data rates of 1.7048 and 1.756, respectively; see step 5 above). This is shown by striking out the row and column pertaining to link l_1 in Fig. 2(e)

(step 21) Transmitter C removes link l_1 from $ActiveSet$ (since link l_1 , having been eliminated in this round, is no longer in contention with the other links in $ActiveSet$ for carrying transmission in the current time slot). Thus, $ActiveSet = \{l_2, l_3, l_4\}$ is the revised set of links still contending

Algorithm 1

```

1: procedure LC-ELDR(ActiveSet)
2:    $I \leftarrow$  all links in the network for which this node17 is the
   transmitter
3:   while  $|ActiveSet| > 2$  do
4:     Compute  $\tau_{l,max} \forall l \in ActiveSet$  from the delay table after
     suppressing the rows and columns corresponding to links
     not in ActiveSet
5:     Compute the queue-length weighted expected data rate
     that will be realized if link  $l$  is allowed to carry trans-
     mission given the  $\tau_{l,max}$ -delayed channel state of link  $l$ ,
     for all links  $l$  in ActiveSet, as follows:  $Q_l[t - \tau_{l,max}] \times$ 
      $E[C_l[t] | C_l[t - \tau_{l,max}] = c_{l,\tau_{l,max}}]$ , where  $c_{l,\tau_{l,max}}$  is
     the channel-state realization of link  $l$  at time  $t - \tau_{l,max}$ 
6:     Let  $H$  be the link with the largest queue-length weighted
     expected data rate computed in Step 5 (arbitrarily chosen
     if more than one satisfy this criterion)
7:     Let EC (for Elimination Candidates) be the subset of
      $ActiveSet \setminus \{H\}$  such that if link  $K \in ActiveSet \setminus \{H\}$ 
     then link  $K \in EC$  if, on recomputing the delays
      $\tau_{l,max} \forall l \in ActiveSet$  after masking the rows and
     columns corresponding to link  $K$  and the links not
     in ActiveSet from the table of delay values, there is
     a reduction in the  $\tau_{l,max}$  value of at least one link
      $l \in ActiveSet \setminus \{K\}$ 
8:     if  $EC = \emptyset$  (the empty set) then
9:       if  $H \in I$  then
10:        Set transmit decision of  $H = TRANSMIT$ ; For
        all  $l \in I \setminus \{H\}$ , set transmit decision of  $l =$ 
        NOTTRANSMIT
11:      else
12:        For all  $l \in I$ , set transmit decision of  $l =$ 
        NOTTRANSMIT
13:      end if
14:      Exit.
15:    else
16:      Let  $S$  be the link in EC with the lowest queue-
      length weighted expected data rate computed in Step
      5 (arbitrarily chosen if more than one satisfy this
      criterion)
17:      if  $S \in I$  then
18:        Set transmit decision of  $S =$  NOTTRANSMIT.
19:         $I \leftarrow I \setminus \{S\}$ 
20:        if  $I = \emptyset$  then Exit.
21:        end if
22:         $ActiveSet \leftarrow ActiveSet \setminus \{S\}$ 
23:      end if
24:    end while
25:    Recompute the delays  $\tau_{l,max}$  as in Step 4. Recompute the
     queue-length weighted expected data rates as in Step 5. Let
      $T \in ActiveSet$  be the link with the largest expected data rate
     (arbitrarily chosen if more than one satisfy this criterion)
26:    if  $T \in I$  then
27:      Set transmit decision of  $T = TRANSMIT$ . For all  $l \in$ 
       $I \setminus \{T\}$ , set transmit decision of  $l =$  NOTTRANSMIT
28:    else
29:      For all  $l \in I$ , set transmit decision of  $l =$  NOTTRANSMIT
30:    end if
end procedure

```

ROUND 2:

Note: For all purposes, the table of delay values for this round is the original table minus the row and column corresponding to link l_1 since it was eliminated in the previous round.

(step 4) transmitter C sets $\tau_{l_1,max} = X$ (don't care) since link l_1 does not belong to *ActiveSet* anymore. Transmitter C then calculates the new values of $\tau_{l_i,max}$, for $i = 2, 3, 4$. All these values could potentially be smaller than the corresponding values in round 1, since when calculating these values in round 2, transmitter C ignores the delay values in the row and column pertaining to link l_1 which was eliminated in the previous round. The new values are: $\tau_{l_2,max} = 2$, $\tau_{l_3,max} = 5$, $\tau_{l_4,max} = 1$. Note that $\tau_{l_4,max}$ has reduced from 3 (in the previous round) to 1 (see Fig. 2(e)).

(step 5) as in round 1, transmitter C computes the expected data rate on each link l_i in *ActiveSet* given the $\tau_{l_i,max}$ -delayed channel-state realization of that link (as computed in step 4 above), as follows: $E[C_{l_2}[t] | C_{l_2}[t-2] = 2] = 1.82$, $E[C_{l_3}[t] | C_{l_3}[t-5] = 2] = 1.6638$, $E[C_{l_4}[t] | C_{l_4}[t-1] = 1] = 1$.

(step 6) as in round 1, transmitter C sets aside the link that has the largest expected data rate in this round, as computed in the previous step. Coincidentally, link l_2 turns out to be that link in this round also. Therefore, transmitter C sets $H = l_2$. This is shown by highlighting transmitter B 's column in green in Fig. 2(f).

(step 7) once again, transmitter C goes about computing *EC* (to decide on the link it wants to eliminate in this round) as follows:

- first, transmitter C decides to skip this procedure for link l_1 (since it is not in *ActiveSet* anymore), and for link l_2 (since l_2 , being the link with the largest expected data rate, was set aside in step 6 above)
- next, following similar arguments as in round 1, transmitter C decides that link l_3 (i.e., its own link) is not a candidate for elimination (see Fig. 2(f))
- lastly, again following similar arguments as in round 1, transmitter C decides that link l_4 is a candidate for elimination, and adds it to the set of elimination candidates, *EC* (see Fig. 2(g)).

Thus, transmitter C has computed the set of elimination candidates to be $EC = \{l_4\}$

(step 16) since there is only one elimination candidate (namely, link l_4) in set *EC*, transmitter C chooses to eliminate link l_4 by setting $S = l_4$. This is shown by striking out the row and column pertaining to link l_4 in Fig. 2(h)

(step 21) Transmitter C removes link l_4 from *ActiveSet* (since link l_4 , having been eliminated in this round, is no longer in contention with the other links in *ActiveSet* for carrying transmission in the current time slot). Thus, $ActiveSet = \{l_2, l_3\}$ is the revised set of links still contending

ROUND 3:

(step 24) transmitter C sets $\tau_{l_1,max}$ and $\tau_{l_4,max}$ to X (don't care) since links l_1 and l_4 do not belong to *ActiveSet* anymore. Transmitter C then calculates the new values of $\tau_{l_i,max}$,

for $i = 2, 3$ to be $\tau_{l_2, \max} = 1$, $\tau_{l_3, \max} = 1$. Transmitter C computes the expected data rate on each link l_i in $ActiveSet$ given the $\tau_{l_i, \max}$ -delayed channel-state realization of that link (where the new $\tau_{l_i, \max}$ value is as computed only just), as follows: $E[C_{l_2}[t] | C_{l_2}[t-1] = 1] = 1.1$, $E[C_{l_3}[t] | C_{l_3}[t-1] = 2] = 1.9$. Note that $\tau_{l_2, \max}$ and $\tau_{l_3, \max}$ have changed from 2 and 5 respectively (in the previous round) to 1 each (see Fig. 2(h)). Transmitter C sets $T = l_3$ since link l_3 (i.e., its own link) has the largest expected data rate as computed above

(step 26) Transmitter C sets the transmit decision for link l_3 (i.e., for its own link) to TRANSMIT

(step 30) Transmitter C terminates this run of the $LC-ELDR$ policy.

It is easy to convince oneself that all the transmitters running the $LC-ELDR$ policy in the current time slot would arrive at the same decision – links other than l_3 would set their transmit decision to NOTTRANSMIT while link l_3 sets its transmit decision to TRANSMIT. To illustrate this in short, transmitter D for example, when running the $LC-ELDR$ policy in the current time slot, would make the same calculations as shown above till step 16 in round 2, whereupon transmitter D would execute step 18 (note that $I = \{l_4\}$ in this case) and thereby set the transmit decision for link l_4 (i.e., its own link) to NOTTRANSMIT, and terminate this run of the algorithm at transmitter D at step 19.

So far, we have considered the case of networks with complete interference. Extension to the case of networks with multiple interference sets is as follows: (i) Let $ActiveSet$ contain all the links in the network. Set $M \leftarrow \phi$. Set $\delta_{ij} \leftarrow 0 \forall l_i \notin I_{l_j} \forall i, j$ where δ_{ij} is the delay value in row i and column j of the delay table (ii) Run the $LC-ELDR$ policy. Let l be the link chosen for transmission by the $LC-ELDR$ policy. Set $M \leftarrow M \cup \{l\}$. Set $ActiveSet \leftarrow ActiveSet \setminus (\{l\} \cup I_l)$. (iii) Repeat step (ii) while $ActiveSet$ is non-empty. When $ActiveSet$ becomes empty, the links in M are the set of links that will be allowed to carry transmission in the current time slot.

The policy $LC-ERDMC$ is a minor variation of $LC-ELDR$, where only step 16 of $LC-ELDR$ is modified as noted in Algorithm 2. In the example above, in step 16 of round 1, the $LC-ELDR$ policy chose to eliminate link l_1 since it had a lower expected data rate compared to that of link l_4 , whereas the $LC-ERDMC$ policy chooses to eliminate link l_4 since eliminating link l_4 reduces the delays with which the channel state of two channels (namely, l_2 and l_3) can be accessed, whereas eliminating link l_1 would reduce the delay of only one channel (namely, l_4).

B. Throughput Near-Optimality

In this section, we derive analytical expressions for the expected saturated system throughputs of the $LC-ELDR$ and $LC-ERDMC$ policies. We evaluate these expressions in Sec. VII for various network settings, and demonstrate that these expressions approximate the optimal throughput values very closely. We need a few definitions, which we introduce here rather informally, and make these definitions precise

Algorithm 2

procedure $LC-ERDMC(ActiveSet)$

16: Let S be the link in EC with the lowest expected data rate among those links that, upon their elimination, reduce the maximum of the delay values in a row, for the largest number of rows (i.e., channels) [corresponding to links in $ActiveSet$] in the table of delay values [after suppressing the rows and columns corresponding to links that are not in $ActiveSet$]

end procedure

mathematically later in Appendices L and M. Let N be the number of links and $T := \{1, 2, \dots, N\}$ be the set of links in $ActiveSet$ before calling the $LC-ELDR$ policy. Further, let

$$\mathbf{1}\{a\} := \begin{cases} 1 & \text{if } a \text{ is true} \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbf{1}_+\{a_i\}_{i \in I} := \max \{\mathbf{1}\{a_i\}\}_{i \in I}$$

We will call the “while” loop body from line 3 to line 23, and also the computation in lines 24–29 in the algorithm listing of the $LC-ELDR$ policy in Sec. VI, a “round”. Thus, if the $LC-ELDR$ policy terminates at line 26 or at line 28, then it would have executed $N - 1$ rounds (specifically, $N - 2$ rounds in the body of the “while” loop, and the last round [round $N - 1$] in lines 24–29), and it would have executed r ($r < N - 1$) rounds if it terminates at line 14 or 19. Thus, the $LC-ELDR$ policy can terminate after r rounds, $1 \leq r \leq N - 1$.

Let $\tau_j^{(r)}$ (not to be confused with $\tau_l(h)$ defined in Sec. II-D) be equal to $\tau_{l_j, \max}$ (see Sec. II-D) at the beginning of round r , where $\tau_{l_j, \max}$ is calculated after masking the rows and columns pertaining to the links that have been eliminated in rounds 1 to $r - 1$ (as illustrated in Figs. 2(a), 2(e), and 2(h)). Let $c_{k, \tau}$ be the realization of channel-state on link l_k at time $t - \tau$.

$e^{(r)}$: Let $e^{(r)}$ be the link that was (or that will be) eliminated (by the $LC-ELDR$ policy) in round r (thus, “ $e^{(r)} = k$ ” [for $k \in T$] implies that link k was (or will be) eliminated in round r , and “ $e^{(r)} = 0$ ” implies that there is no candidate link to eliminate, and hence that the $LC-ELDR$ policy terminates).

Let $i \in T$ be the link chosen by the $LC-ELDR$ policy in a particular time slot, say t , when the $LC-ELDR$ policy terminates after executing r rounds, $1 \leq r \leq N - 1$. We now consider the working of the $LC-ELDR$ policy when it is executed at the transmitter node of link i in time slot t . Consider a particular round \tilde{r} ($1 \leq \tilde{r} \leq r$) of the $LC-ELDR$ policy. If link i has survived (i.e., was not eliminated in) the first $\tilde{r} - 1$ rounds, then it will survive round \tilde{r} if one of the following happens:

- (i) $s(i, \tilde{r}, 1)$: link i has the largest expected data rate in round \tilde{r} (hence $i = H$ and therefore $i \notin EC$; see steps 6 and 7 of Algorithm 1). We denote this condition as $s(i, \tilde{r}, 1)$.
- (ii) $s(i, \tilde{r}, 2)$: link i does not have the largest expected data rate in round \tilde{r} (hence $i \neq H$), and link i does not reduce the maximum of the delay values in a row, for any row (corresponding to a link in $ActiveSet$) in the table of delay values (after suppressing the rows and columns corresponding to links that have been eliminated up to round \tilde{r}) if it is eliminated in round \tilde{r} (hence $i \notin EC$;

see step 7 of Algorithm 1), and some link (other than the one with the largest expected data rate in round \tilde{r}) reduces the maximum of the delay values in a row, for at least one row if it is eliminated in round \tilde{r} (i.e., $EC \neq \phi$). We denote this condition as $s(i, \tilde{r}, 2)$.

(iii) $s(i, \tilde{r}, 3)$: link i does not have the largest expected data rate in round \tilde{r} (hence $i \neq H$), and link i reduces the maximum of the delay values in a row, for at least one row (corresponding to a link in *ActiveSet*) in the table of delay values (after suppressing the rows and columns corresponding to links that have been eliminated up to round \tilde{r}) if it is eliminated in round \tilde{r} (hence $i \in EC$), and link i does not have the smallest expected data rate among the links that reduce the maximum of the delay values in a row, for at least one row if that link is eliminated in round \tilde{r} (hence $i \neq S$; see step 16 of Algorithm 1). We denote this condition as $s(i, \tilde{r}, 3)$.

With the required definitions in place, we are now ready to state the following crucial result:

Proposition 10. *The expected saturated system throughput of the LC-ELDR policy is exactly*

$$\begin{aligned} & \sum_{i=1}^N \sum_{r=1}^{N-1} \sum_{\substack{c_{j, \tau_j^{(q)}} \\ j \in T \\ q \in \{1, 2, \dots, r\}}} \mathbb{E} \left[C_{l_i}[t] \prod_{\tilde{r}=1}^{r-1} \left(\mathbf{1}_{+} \{s(i, \tilde{r}, n)\}_{n \in \{1, 2, 3\}} \right. \right. \\ & \quad \left. \times \mathbf{1}_{\{e^{(\tilde{r})} \neq 0\}} \right) \mathbf{1}_{\{s(i, r, 1)\}} \mathbf{1}_{\{e^{(r)} = 0\}} \\ & \quad \left| C_{l_i}[t - \tau_i^{(r)}] = c_{i, \tau_i^{(r)}}, \right. \\ & \quad \left. \left\{ C_{l_i}[t - \tau_i^{(q)}] = c_{i, \tau_i^{(q)}} \right\}_{q \in \{1, 2, \dots, r-1\}}, \right. \\ & \quad \left. \left\{ C_{l_p}[t - \tau_p^{(q)}] = c_{p, \tau_p^{(q)}} \right\}_{\substack{p \in T \setminus \{i\} \\ q \in \{1, 2, \dots, r\}}} \right] \\ & \quad \times \prod_{m=1}^N \left(\pi_m(c_{m, \tau_m^{(1)}}) \prod_{n=1}^{r-1} p_{m, c_{m, \tau_m^{(n)}}}^{(\tau_m^{(n)} - \tau_m^{(n+1)})} c_{m, \tau_m^{(n+1)}} \right) \end{aligned}$$

where $\pi_m(c)$ is the steady-state probability of being in state c in the channel-state Markov chain of link m , and $p_{m, ij}^{(k)}$ is the transition probability of reaching state j from state i in k steps in the channel-state Markov chain of link m , with $p_{m, ij}^{(0)} = 1 \forall i, j, m$.²⁰

We present mathematically precise definitions of the symbols that were defined informally in the paragraphs preceding Proposition 10, and the proof itself in Appendix L.

Proposition 11. *The expression for the expected saturated system throughput of the LC-ERDMC policy is the same as the expression in Proposition 10, except that the expression for $e^{(r)}$ is redefined to be as noted in Appendix M.*

The proof of this proposition is essentially the same as that for Proposition 10; the minor departures are noted in Appendix M. We evaluate the expressions in Propositions 10 and 11

²⁰ $\pi_m(c) = \pi(c)$ and $p_{m, ij}^{(k)} = p_{ij}^{(k)}$ $\forall m \in \mathcal{L}$ in our model; see Sec. II-A

numerically and plot them for various network settings in Sec. VII, and demonstrate that these expressions approximate the optimal throughput values very closely and hence that the LC-ELDR and LC-ERDMC policies are near-throughput-optimal.

C. Computational Complexity

We have the following result:

Proposition 12. *The running times of both of the LC-ELDR and LC-ERDMC policies are $\mathcal{O}(CL^2 + L^4)$.*

We provide a proof of this result in Appendix K. Thus from this proposition, for the case of a network with complete interference, the computational complexities of the LC-ELDR and LC-ERDMC policies are polynomial in the number links in the network and linear in the number of channel states. For a network with multiple interference sets, in the worst case (which happens when $I_l = \phi$, $\forall l \in \mathcal{L}$), there are L calls to Algorithm 1, thus there is an additional multiplicative factor of L , which does not alter the (orders of) polynomial and linear dependencies on the number of links and the number of channel states respectively. Thus, the LC-ELDR and LC-ERDMC policies offer an enormous improvement over the computational complexities of the R and H policies. We compare the running times of these policies in Table VII.

VII. NUMERICAL RESULTS AND DISCUSSION

In this section, we present numerical results comparing the saturated system throughputs of the R , H , LC-ELDR and LC-ERDMC policies. In addition, for purposes of comparison, we define two new policies that we call O and IC :

Policy O : $\arg \max \{Q_l[t - \tau_{l, \max}] \times \mathbb{E}[C_l[t] | C_l[t - \tau_{l, \max}] = c_{l, \tau_{l, \max}}]\}_{l \in \mathcal{L}}$.

Policy IC : $\arg \max \{Q_l[t] \times C_l[t]\}_{l \in \mathcal{L}}$.

The O policy chooses the link l that has the largest value of queue-length weighted expected data rate on the link, computed using the $\tau_{l, \max}$ -delayed NSI of that link. Note that the O policy corresponds to stopping after round 1 in the LC-ELDR and LC-ERDMC policies and declaring link H the “winner” (see step 6 of Algorithm 1). The IC policy chooses the link l that has the largest value of queue-length weighted data rate on the link, computed using instantaneous NSI of that link, with the assumption that each link has access to the instantaneous NSI for all channels. We emphasize that the IC policy does not conform to the structure of delayed NSI noted in Sec. II-D, but it serves as an upper bound on the system throughput that can be achieved in the delayed NSI regime.

A. Methodology

As in [1], we consider networks with complete interference (i.e., $I_l = \mathcal{L} \setminus \{l\}$, $\forall l \in \mathcal{L}$) and perfect collision (i.e., $\gamma_l = 0$, $\forall l \in \mathcal{L}$). The channel state on each link is modeled as

TABLE IV: The *Very Small Delays* (VSD) delay profile for a wireless network with three links

	At TX A	At TX B	At TX C
Delay in obtaining NSI of link l_1	0	1	1
Delay in obtaining NSI of link l_2	1	0	1
Delay in obtaining NSI of link l_3	1	2	0

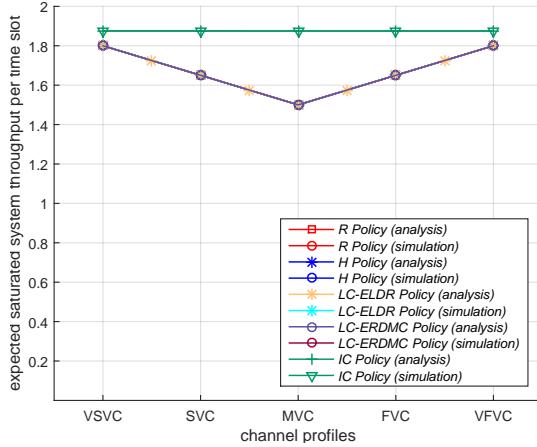


Fig. 3: Expected saturated system throughputs (in data units transmitted) per time slot of the various policies for a network with three links, VSD delay profile and different channel profiles. The curves of all but the IC policy overlap.

a DTMC on the state space $\mathcal{C} = \{1, 2\}$. We only consider single-hop transmissions in the network. We implemented all the policies in a custom C++ simulator. All the numbers from simulation that we quote in this section were averaged over 10^7 trials.

B. Results and Discussion - Throughput

Fig. 3 shows the expected saturated system throughputs of the various policies for a network with three links, for the *very small delay* (VSD) delay profile (see Table IV), for different channel profiles. For this purpose, we define five channel profiles – namely, *very slow varying channel* (VSVC), *slow varying channel* (SVC), *medium varying channel* (MVC), *fast varying channel* (FVC) and *very fast varying channel* (VFVC), with channel-state crossover probabilities of 0.1, 0.3, 0.5, 0.7 and 0.9 respectively. We observe that all the policies perform equally well when the delay values are very small (the plots for all but the IC policy overlap in Fig. 3).

Fig. 4 shows the expected saturated system throughputs of the *LC-ELDR*, *LC-ERDMC*, *O* and *IC* polices for a network with three links and VSVC channel profile, for different delay profiles. For this purpose, in addition to the VSD delay profile noted in Table IV, we define four other delay profiles – namely, *small delays* (SD), *medium delays* (MD), *large delays* (LD), and *very large delays* (VLD) as shown below:

$$SD : \begin{bmatrix} 0 & 1 & 3 \\ 2 & 0 & 4 \\ 1 & 2 & 0 \end{bmatrix}, \quad MD : \begin{bmatrix} 0 & 7 & 11 \\ 8 & 0 & 9 \\ 12 & 6 & 0 \end{bmatrix},$$

$$LD : \begin{bmatrix} 0 & 20 & 15 \\ 24 & 0 & 17 \\ 12 & 28 & 0 \end{bmatrix}, \quad VLD : \begin{bmatrix} 0 & 78 & 36 \\ 59 & 0 & 88 \\ 45 & 92 & 0 \end{bmatrix}$$

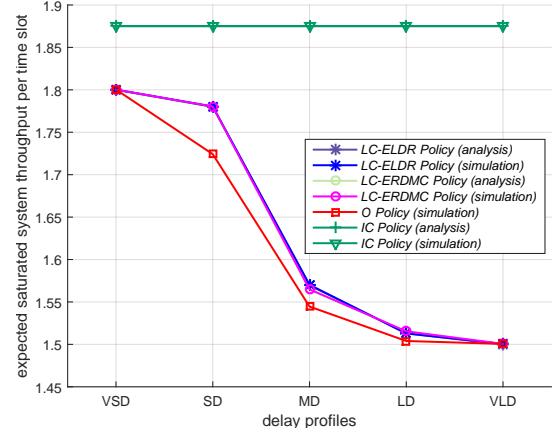


Fig. 4: Expected saturated system throughputs (in data units transmitted) per time slot of *LC-ELDR*, *LC-ERDMC*, *O* and *IC* policies for a network with three links, VSVC channel profile and different delay profiles. The analysis and simulation curves of each policy overlap. Throughput of *LC-ERDMC* is marginally below and marginally above that of *LC-ELDR* for the MD and LD delay profiles respectively.

TABLE V: Table of heterogeneous delay values for the *Very Small Delays* (VSD) delay profile for a network with ten links

	At TX T_1	At TX T_2	At TX T_3	At TX T_4	At TX T_5	At TX T_6	At TX T_7	At TX T_8	At TX T_9	At TX T_{10}
Link l_1	0	1	1	1	1	1	1	1	1	1
Link l_2	1	0	1	1	1	1	1	1	1	1
Link l_3	1	2	0	1	1	1	1	1	2	1
Link l_4	2	1	1	0	1	1	1	2	1	1
Link l_5	1	1	1	1	0	1	2	1	1	2
Link l_6	1	2	1	1	1	0	1	1	1	1
Link l_7	1	1	1	1	1	1	0	2	1	1
Link l_8	1	2	1	1	2	1	1	0	2	1
Link l_9	2	1	1	1	1	1	1	2	0	1
Link l_{10}	1	1	2	1	1	1	2	1	1	0

Similar to the entries in Table IV, the entry at position ij in each of the delay profiles is to be interpreted as the least delay with which the NSI of link l_i is available at transmitter j (i.e., at the transmitter node of link l_j). We do not show the system throughputs of the *R* and *H* policies in Fig. 4 since the delay values for all but the VSD profile are such that it is computationally hard to evaluate the saturated system throughputs of these policies. We note that the *LC-ELDR* and *LC-ERDMC* policies perform almost equally well, and both of them outperform the *O* policy. Fig. 4 also depicts the degradation of system throughput with increasing delay values.

Fig. 5 shows the expected saturated system throughputs of the *R*, *H*, *LC-ELDR*, *LC-ERDMC*, *O* and *IC* polices for the case of VSVC channel profile and VSD delay profile, for different number of links in the network. The VSD delay profile, noted in Table V, is for a network with ten links. The VSD delay profile for a network with n links, $2 \leq n \leq 10$, is derived from this table by taking the $n \times n$ sub-matrix hinged at the top-left corner of this table. We show analysis plots of the *LC-ELDR* and *LC-ERDMC* policies only up to $\#links = 5$ since computing these values for networks with larger number of links is computationally costly. We note that the throughputs of both the *LC-ELDR* and *LC-ERDMC* policies overlap with the optimal throughput values of the *R* and *H* policies.

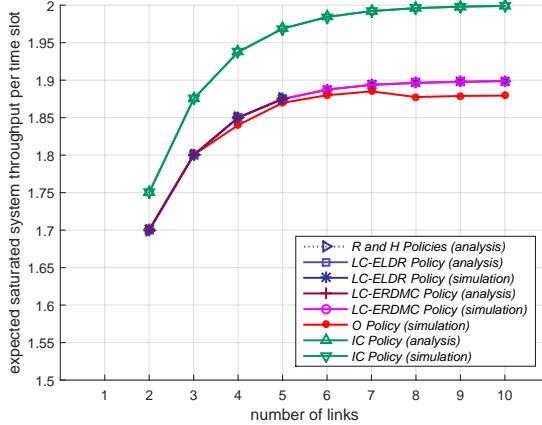


Fig. 5: Expected saturated system throughputs (in data units transmitted) per time slot of the R , H , LC - $ELDR$, LC - $ERDMC$, O and IC policies for a network with VSVC channel profile and VSD delay profile, for different number of links in the network (analysis curves of LC - $ELDR$ and LC - $ERDMC$ are shown only up to #links = 5). The analysis and simulation curves of each policy overlap.

TABLE VI: Table of heterogeneous delay values for the *Medium Delays* (MD) delay profile for a network with ten links

At	At	At	At	At	At	At	At	At	At	
TX	TX	TX	TX	TX	TX	TX	TX	TX	TX	
T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9	T_{10}	
Link l_1	0	7	11	6	5	3	9	7	11	4
Link l_2	8	0	9	7	3	2	5	9	2	8
Link l_3	12	6	0	11	3	4	11	7	2	9
Link l_4	9	11	2	0	5	7	2	1	10	5
Link l_5	2	5	11	3	0	7	8	9	10	4
Link l_6	1	9	2	4	8	0	11	6	5	2
Link l_7	12	1	3	5	9	11	0	7	9	1
Link l_8	7	7	1	2	11	8	4	0	11	8
Link l_9	4	1	4	4	9	12	11	7	0	1
Link l_{10}	1	1	12	4	7	1	1	9	12	0

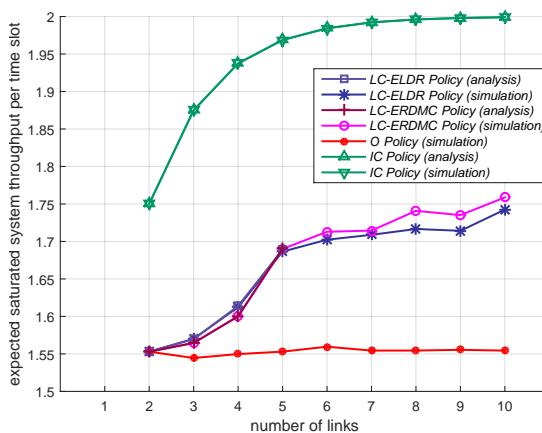


Fig. 6: Expected saturated system throughputs (in data units transmitted) per time slot of the LC - $ELDR$, LC - $ERDMC$, O and IC policies for a network with VSVC channel profile and MD delay profile, for different number of links in the network (analysis curves of LC - $ELDR$ and LC - $ERDMC$ are shown only up to #links = 5). The analysis and simulation curves of each policy overlap.

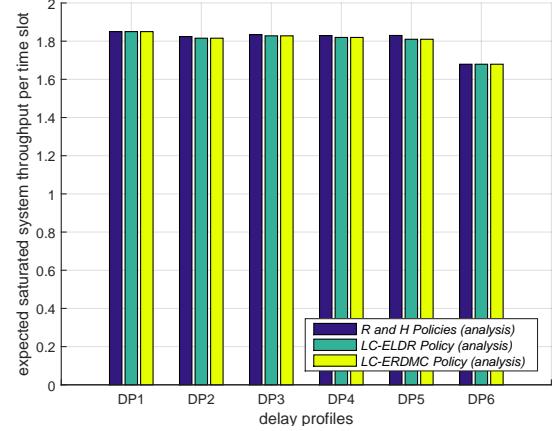


Fig. 7: Expected saturated system throughputs (in data units transmitted) per time slot of the R , H , LC - $ELDR$, LC - $ERDMC$ policies for networks with four links, VSVC channel profile and different delay profiles (DP).

Fig. 6 shows the expected saturated system throughputs of the LC - $ELDR$, LC - $ERDMC$, O and IC polices for the case of VSVC channel profile and MD delay profile, for different number of links in the network. The MD delay profile for a network with ten links is shown in Table VI. As before, the MD delay profile for a network with n links, $2 \leq n \leq 10$, is derived from this table by taking the $n \times n$ sub-matrix hinged at the top-left corner of this table. We do not show the system throughputs of the R and H polices in Fig. 6 since the MD delay profile is such that it is computationally hard to evaluate the saturated system throughputs of these policies. Again, as before, we show analysis plots of the LC - $ELDR$ and LC - $ERDMC$ policies only up to #links = 5 since computing these values for networks with more links is computationally costly. We note that both of the LC - $ELDR$ and LC - $ERDMC$ policies outperform the O policy. We also note, more importantly, that there are situations (#links = 3, 4) where LC - $ELDR$ outperforms LC - $ERDMC$ and other situations (#links = 6 to 10) where LC - $ERDMC$ outperforms LC - $ELDR$, implying that none of these policies could be throughput optimal.²¹

Fig. 7 shows the expected saturated system throughputs of the R , H , LC - $ELDR$ and LC - $ERDMC$ polices for networks with four links, VSVC channel profile and different delay profiles (DP for short; see below). We first comment that the high computational cost of the R and H policies hinder us from measuring their performances for very many delay profiles, and for all the delay profiles that we could measure, the expected saturated system throughputs of the LC - $ELDR$ and LC - $ERDMC$ policies were less than 1.1% (and in many cases less than 0.5%) away from the optimal

²¹With reference to the low-complexity scheduling policies in Sec. VI, in the process of obtaining a link that would maximize the expected system throughput, in each iteration, the LC - $ELDR$ policy discards the link with the lowest expected data rate, and the LC - $ERDMC$ policy discards the link that, upon its elimination, reduces the delay values for the largest number of channels (see step 16 of the algorithms). Even though taking recourse to eliminating one link in each iteration gives the LC - $ELDR$ and LC - $ERDMC$ policies their short running times, the approaches that these policies take to achieve this, as alluded to in Fig. 6, are not optimal. It remains unresolved as to whether there is an optimal strategy to isolate a link that would be eliminated in each iteration, and if one exists, what its structure should be. This needs further exploration.

TABLE VII: Comparison of the running times of the R , $LC-ELDR$ and $LC-ERDMC$ policies. All times are in microseconds.

# Links	R Policy ^{†‡}	$LC-ELDR$ Policy [‡]	$LC-ERDMC$ Policy [‡]
5	4×10^3	11	10
10	1.1×10^5	102	96
15	3.9×10^6	165	151
20	8.14×10^7	247	237

[†] As reported in [1].

[‡] Information about the delay profile used in [1] is not available. We are reasonably sure that the delay profile used in [1] uses smaller delay values than those in the MD delay profile that we have used here for the $LC-ELDR$ and $LC-ERDMC$ policies, since evaluating the R policy for the MD delay profile is prohibitively time consuming. Thus the running time values that we mention here for the $LC-ELDR$ and $LC-ERDMC$ policies are relatively pessimistic, and provide a more-than-fair comparison.

[‡] As measured on a computer with Intel Core 2 Duo processor clocked at 2 GHz, with 2 GB RAM.

[‡] For the MD delay profile whose delay values for #links = 5, 10 are shown in Table VI. The delay values in the MD delay profile for #links = 15, 20 are similar to those shown in Table VI.

throughput values of the R and H policies, thus demonstrating the near-optimality of throughputs of the $LC-ELDR$ and $LC-ERDMC$ policies. The following are the definitions of the six delay profiles referred to in Fig. 7 (the entry at position ij has the same interpretation as before):

$$DP1 : \begin{bmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix}, \quad DP2 : \begin{bmatrix} 0 & 1 & 1 & 2 \\ 2 & 0 & 2 & 2 \\ 2 & 2 & 0 & 1 \\ 2 & 2 & 2 & 0 \end{bmatrix},$$

$$DP3 : \begin{bmatrix} 0 & 1 & 1 & 2 \\ 2 & 0 & 2 & 2 \\ 2 & 2 & 0 & 1 \\ 2 & 1 & 1 & 0 \end{bmatrix}, \quad DP4 : \begin{bmatrix} 0 & 2 & 2 & 3 \\ 1 & 0 & 2 & 2 \\ 2 & 2 & 0 & 2 \\ 1 & 2 & 1 & 0 \end{bmatrix},$$

$$DP5 : \begin{bmatrix} 0 & 2 & 2 & 2 \\ 1 & 0 & 2 & 2 \\ 2 & 3 & 0 & 3 \\ 1 & 2 & 1 & 0 \end{bmatrix}, \quad DP6 : \begin{bmatrix} 0 & 4 & 4 & 4 \\ 4 & 0 & 4 & 4 \\ 4 & 4 & 0 & 4 \\ 4 & 4 & 4 & 0 \end{bmatrix}$$

In Table VII, we compare the running times of the R , $LC-ELDR$ and $LC-ERDMC$ policies for a network with varying number of links, VSVC channel profile, and MD delay profile. *Evidently, in stark contrast to the R policy, the $LC-ELDR$ and $LC-ERDMC$ policies are computationally very efficient, validating the computational complexity analyses in Propositions 3, 4 and 12.*

C. Results and Discussion – Queueing Delay

In this section, we evaluate the queueing delay performances of the $DQIC1$, $DQIC2$, R and H policies numerically and compare these with the queueing delay values obtained by evaluating the analytical expression in Eq. (16) for these policies. In Sec. V-A, we saw that the computational cost of comprehensively evaluating Expr. (16) is prohibitively expensive even for a small network that has only 2 transmitters, when there are only two channel states for any link, and when the number of packet arrivals in any slot is only 0 or 1, for a small number of time slots $T = 10$. While reducing T to smaller values (say $T = 5$) will allow for a comprehensive evaluation of Expr. (16), small values of T do not allow us to capture the dynamics of packet arrivals (and hence of queue lengths) and of channel state transitions precisely enough. In fact, the number of time slots (T) over which we evaluate Expr.

(16) should be at least of the order of a few hundred to be able to capture the queue and channel transition dynamics with reasonable accuracy. Therefore, we resort to evaluating Expr. (16) not comprehensively but for a few arrival streams and for a few channel state sample paths. We do this in two ways – (i) by selecting these sample arrival streams and channel state sample paths randomly, and (ii) by selecting typical (i.e., high probability) sample arrival streams and typical channel state sample paths. This gives us leeway in evaluating Expr. (16) for larger values of T .

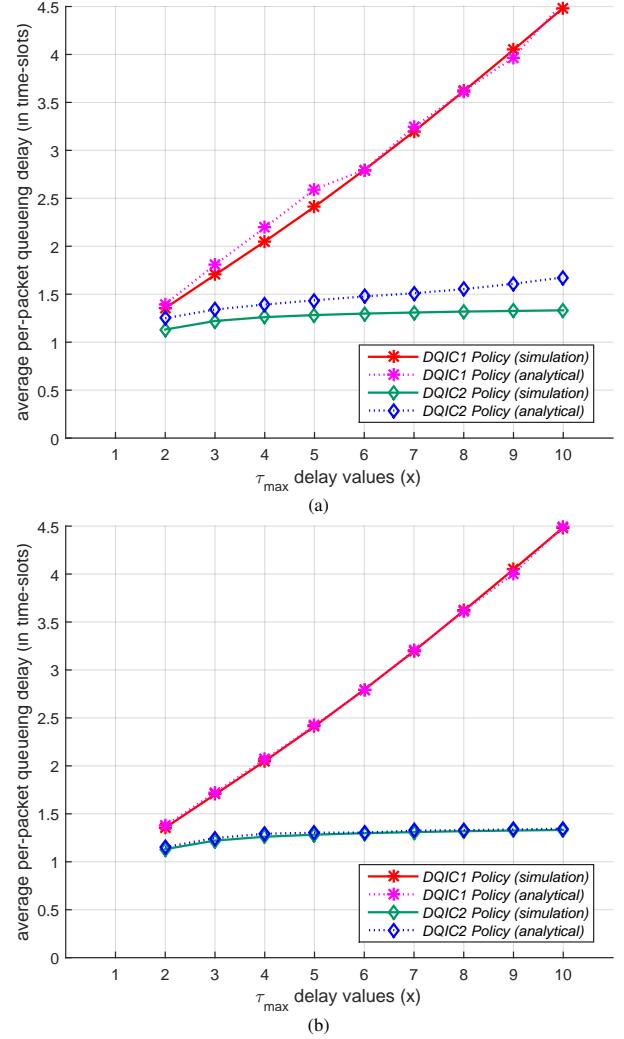


Fig. 8: Comparison of analytical and simulation results for the average queueing delay per-packet in the $DQIC1$ and $DQIC2$ policies. The $DQIC1$ policy uses τ_{max} -delayed QSI whereas the $DQIC2$ policy uses τ_{max} -delayed QSI. The analytical queueing delay values are obtained by using 100 randomly chosen arrival streams and channel state sample paths in (a) and using 100 typical (i.e., high probability) arrival streams and channel state sample paths in (b). The match between analytical and simulation results is substantially more precise when using typical sequences for arrival streams and channel transitions.

First, as in Sec. III-C, we consider a network with two transmitters (links), with the channels on all links modeled as independent Markov chains on the state space $\{1, 2\}$ with channel state transition probability 0.1. We consider heterogeneous delays as noted in Table II, where $x \geq 1$ is a parameter we vary. Packets arrive into the two queues at

the two transmitters as independent Poisson processes with rates $\lambda_1 = \lambda_2 = \frac{1}{4}$ (for this setting, it is easily seen that sum rates (i.e. $\lambda_1 + \lambda_2$) of up to 1.75 are supportable; see footnote 15). In Fig. 8, we compare the average queueing delay performances of the *DQIC1* and *DQIC2* policies. For the analytical plots we take the time horizon $T = 2000$ and for the simulation plots we average the delay values over 10^7 trials. For reasons mentioned in the previous paragraph, for the analytical plot in Fig. 8(a), at each of the two links in the network, we consider a 100 random arrival streams and a 100 random channel state sample paths, and for the analytical plot in Fig. 8(b), we consider a 100 typical arrival streams and 100 typical channel state sample paths (this gives a total of 10^4 or 10^8 different arrival stream - channel sample path pairs). We observe that the match between analytical and simulation results is very precise when using typical sequences for arrival streams and channel transitions.

Next, in Fig. 9(a), we compare the average queueing delay performances of the *R* and *H* policies, and in Fig. 9(b), that of *R*, *LC-ELDR* and *LC-ERDMC* policies.²² Because of the prohibitively large computational complexity involved in evaluating the *R* and *H* policies, we consider only 10 typical arrival streams and 10 typical channel state sample paths (for a total of 10^4 different arrival stream - channel sample path pairs). Again, the match between analytical and simulation results is precise validating the gains in delay performance that can be obtained by using τ_{\max} -delayed QSI.

Finally, in Fig. 9(c), concentrating on the delayed queue-lengths at link l_1 , we compare the correlation coefficients between the queue-lengths at $Q_{l_1}[t]$ (i.e., the instantaneous queue-length) and $Q_{l_1}[t-x]$ for $x = 1, \dots, 10$ in the *DQIC1* and *DQIC2* policies. Since in our example, for the case $\tau_{\max} = 2$, the *DQIC2* policy uses the queue-length at time $t-1$ and the *DQIC1* policy uses the queue-length at time $t-2$, we have highlighted these values with small filled circles. As we would expect, we see that the queue-length at time $t-1$ (i.e. at time $t - \tau_{l_1, \max}$) used by the *DQIC2* policy is more positively correlated with the instantaneous queue-length than is the queue-length at time $t-x$ for $x = 2, \dots, 10$ used by the *DQIC1* policy, thus explaining the delay performance gains of the *DQIC2* policy over that of the *DQIC1* policy.

VIII. CONCLUSION

In this work, addressing the problem of distributed scheduling in wireless networks with heterogeneously delayed NSI, we proposed, analyzed and evaluated the performances of two fast and near-throughput-optimal scheduling policies. Enroute, we identified and dealt with two deficiencies (namely, non-optimal delay performance and high computational complexity) in an earlier work in [1]. We showed that the information afforded by the system model could be exploited more aggressively, and that by doing so, immense reduction in computational complexity and substantial reduction in the expected per-packet queueing delay could be obtained. We proposed a provably throughput-optimal scheduling policy that

²²The *LC-ELDR* and *LC-ELDR* are identical for the case of networks with only two links.

embodies these ideas. The proposed fast and near-throughput-optimal scheduling policies champion these ideas further, possess desirable queueing delay characteristics, and have running times that are of the order of microseconds, sufficiently small for use in practical deployments.

REFERENCES

- [1] A. Reddy, S. Banerjee, A. Gopalan, S. Shakkottai, and L. Ying, "On Distributed Scheduling with Heterogeneously Delayed Network-State Information," *Queueing Systems*, vol. 72, pp. 193 – 218, 2012
- [2] L. Tassiulas, and A. Ephremides, "Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multi-hop Radio Networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936 – 1948, 1992
- [3] S. Mohari, and S. Shakkottai, "MaxWeight vs. BackPressure: Routing and Scheduling in Multi-Channel Relay Networks," *IEEE INFOCOM 2013*, pp. 1537 – 1545, 2013
- [4] L. Tassiulas, and A. Ephremides, "Dynamic Server Allocation to Parallel Queues with Randomly Varying Connectivity," *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 466 – 478, 1993
- [5] S. Shakkottai, and A. Stolyar, "Scheduling for Multiple Flows Sharing a Time-Varying Channel: The Exponential Rule," *American Mathematical Society Translations, Series*, vol. 2, pp. 2002, 2000
- [6] C. Joo, and N. Shroff, "Performance of Random Access Scheduling Schemes in Multi-hop Wireless Networks," *Fortieth Asilomar Conference on Signals, Systems and Computers (ACSSC '06)*, pp. 1937 – 1941, 2006
- [7] S. Rajagopalan, D. Shah, and J. Shin, "Network Adiabatic Theorem: An Efficient Randomized Protocol for Contention Resolution," *SIGMETRICS/Performance*, pp. 133 – 144, 2009
- [8] C. Bordenave, D. McDonald, and A. Proutiere, "Performance of Random Medium Access Control An Asymptotic Approach," *SIGMETRICS Perform. Eval. Rev.*, vol. 36, no. 1, pp. 1 – 12, 2008
- [9] P. Marbach, A. Eryilmaz, and A. Ozdaglar, "Achievable Rate Region of CSMA Schedulers in Wireless Networks with Primary Interference Constraints," *IEEE Conference on Decision and Control*, pp. 1156 – 1161, 2007
- [10] P. Gupta, and A. L. Stolyar, "Optimal Throughput Allocation in General Random-Access Networks," *40th Annual Conference on Information Sciences and Systems*, pp. 1254 – 1259, 2006
- [11] J. Liu, and A. L. Stolyar, "Distributed Queue-Length based Algorithms for Optimal End-to-End Throughput Allocation and Stability in Multi-hop Random Access Networks," *45th Allerton Conference on Communication, Control, and Computing, Urbana-Champaign, IL*, 2007
- [12] A. L. Stolyar, "Dynamic Distributed Scheduling in Random Access Networks," *Journal of Applied Probability*, vol. 45, no. 2, pp. 297 – 313, 2008
- [13] L. Jiang, and J. Walrand, "A Distributed CSMA Algorithm for Throughput and Utility Maximization in Wireless Networks," *IEEE/ACM Transactions on Networking*, vol. 18, no. 3, pp. 960 – 972, 2008
- [14] E. Modiano, D. Shah, and G. Zussman, "Maximizing Throughput in Wireless Networks via Gossiping," *ACM SIGMETRICS/Performance*, pp. 27 – 38, 2008
- [15] S. Sanghavi, L. Bui, and R. Srikant, "Distributed Link Scheduling with Constant Overhead," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 35, no. 1, pp. 313 – 324, 2007
- [16] A. Pantelidou, A. Ephremides, and A. Tits, "Joint Scheduling and Routing for Ad-hoc Networks Under Channel State Uncertainty," *WiOpt*, pp. 1 – 8, 2007
- [17] L. Ying, and S. Shakkottai, "On Throughput Optimality with Delayed Network-State Information," *Technical Report*, 2008
- [18] A. Kumar, 2012, "Discrete Event Stochastic Processes: Lecture Notes for an Engineering Curriculum," Retrieved from <http://www.ece.iisc.ernet.in/~anurag/books/anurag/sqpt.pdf>

APPENDIX A PROOF OF PROPOSITION 1

Since $T_{l_i}(\cdot)$ is a function of $CS_{l_i}^R(\mathbf{C}[t](0 : \tau_{\max}))$ and since $\{C_{l_i}[t - \tau_{\max}]\}_{l_i \in \mathcal{L}}$ is made available as part of the computation of the conditional expectation (the *given*) in Eq. (9), the number of parameters to the threshold function $T_{l_i}(\cdot)$ is the number of elements in the set $\mathcal{V}_i := CS_{l_i}^R(\mathbf{C}[t](0 :$

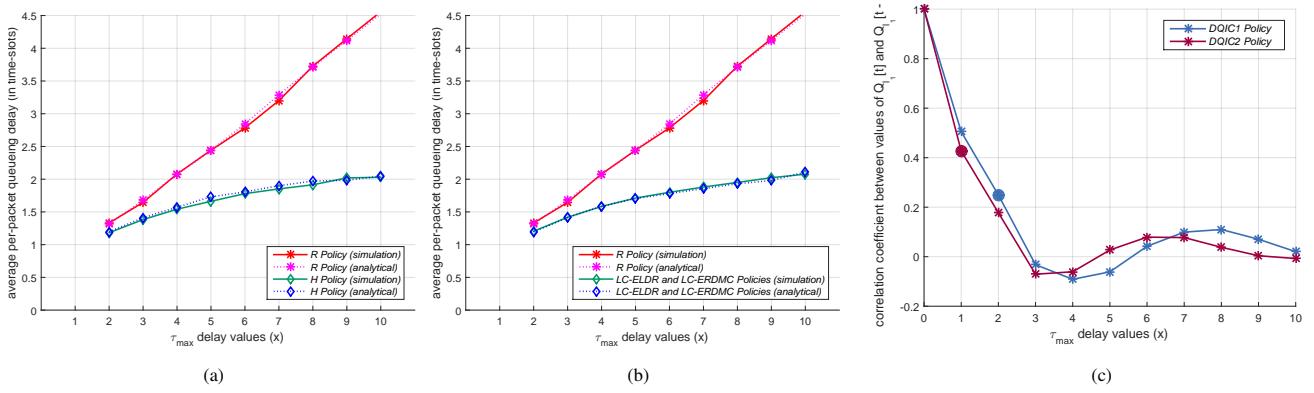


Fig. 9: (a) Comparison of average queueing delay per-packet in the R and H policies. The R policy uses τ_{max} -delayed QSI whereas the H policy uses τ_{max} -delayed QSI. The average per-packet queueing delays of the R policy grows linearly with τ_{max} whereas that of the H policy is sub-linear tending to become almost unaffected by increase in τ_{max} . Comparing the R and H policies and their delay performances, the fact that the R policy uses τ_{max} -delayed queue lengths isolates itself as the sole reason for its undesirable delay performance. (b) Comparison of average queueing delay per-packet in the R and $LC\text{-}ELDR/LC\text{-}ERDMC$ policies. The $LC\text{-}ELDR$ and $LC\text{-}ERDMC$ policies, like the H policy, use τ_{max} -delayed QSI. (c) Comparison of correlation coefficients between values of $Q_{l1}[t]$ and $Q_{l1}[t-x]$ in the $DQIC1$ and $DQIC2$ policies.

$\tau_{max}) \setminus \{C_{l_i}[t - \tau_{max}]\}_{l_i \in \mathcal{L}}$.

Each of the random variables in $CS_{l_i}^R(\cdot)$ can take any of the \mathcal{C} values where \mathcal{C} is the number of channel states. Thus, for $T_{l_i}(\cdot)$, there are $\mathcal{C}^{|\mathcal{V}_i|}$ different values in the domain of $T_{l_i}(\cdot)$. Each of these $\mathcal{C}^{|\mathcal{V}_i|}$ in the domain of $T_{l_i}(\cdot)$ can be independently mapped to a real number. In Proposition 13, we show that it suffices to consider $\mathcal{C} + 1$ carefully chosen real numbers in the range of $T_{l_i}(\cdot)$. For example, from Proposition 13, for $\mathcal{C} = \{1, 2\}$, it suffices to consider the real numbers 0.5, 1.5, 2.5 as the set of possible values in the range of $T_{l_i}(\cdot)$ that each element in the domain of $T_{l_i}(\cdot)$ can be mapped to. Thus, there are $(\mathcal{C} + 1)^{|\mathcal{V}_i|}$ possible choices of threshold functions for $T_{l_i}(\cdot)$. Therefore, considering all the links in the network, the total number of threshold functions that need to be considered in the domain of optimization in Expression (8) is $(\mathcal{C} + 1)^{\sum_{l_i \in \mathcal{L}} \mathcal{C}^{|\mathcal{V}_i|}}$. \square

Proposition 13. *It is sufficient to consider $(\mathcal{C} + 1)$ real numbers in the range of each threshold function $T_l(\cdot)$.*

Proof: Consider the interval $S = (c_i, c_{i+1})$ on the real number line for any integer $i \in \{0, 1, \dots, \mathcal{C}\}$, $c_1, c_2, \dots, c_M \in \mathcal{C}$ (where $c_M = c_C$) are the channel states in the channel state DTMC (see Sec. II-A) and $c_0 = c_1 - 1$.²³ Consider a fixed real number $r_1 \in S$, and a set $T \subseteq \mathcal{C}$. Evidently, r_1 shares the same relationship (i.e., $<$, \leq , $>$, \geq) with each element of T as all other real numbers in S . Hence, it is sufficient to pick r_1 as a representative of all real numbers in S . The result follows noting that there are $\mathcal{C} + 1$ intervals S for each of which a representative real number r_1 has to be chosen. \square

APPENDIX B PROOF OF PROPOSITION 2

$T_l(\cdot)$ is a function of $CS_l^R(\cdot)$, and hence the number of choices of threshold functions for $T_l(\cdot)$ is a function of

²³We use \mathcal{C} to denote both the set and its cardinality. Thus, as set, $\mathcal{C} = \{c_1, c_2, \dots, c_M\}$, and as cardinality, $\mathcal{C} = M$.

the number of random variables in $CS_l^R(\cdot)$, which in turn is dependent on $CS^R(\cdot)$ through Expr. (7). Evidently, the number of random variables in $CS^R(\cdot)$ is maximized when the delays in the table of delay values are all distinct.

The number of random variables in the set $CS_{l_1}^R(\cdot)$ pertaining to each link l_1 (i.e., entries of the form $C_{l_1}[t - \tau]$ in the set $CS_{l_1}^R(\cdot)$) is maximized when $\tau_{l_1}(l_k)$ is the smallest of all the delay values in row k . This is because, the transmitter node of link l_1 , having the CSI of link l_k with the delay $\tau_{l_1}(l_k)$ (which is the smallest delay in row k), also has the CSI of link l_k with delays larger than $\tau_{l_1}(l_k)$, and hence the random variables $\{C_{l_k}[t - \tau]\}_{\tau = \tau_{l_j}(l_k), j=1, \dots, L, j \neq k}$ are all present in the set $CS_{l_1}^R(\cdot)$. Similarly, after fixing the delay value at position $(k, 1)$ to be the smallest value in row k as noted above, the number of random variables in the set $CS_{l_2}^R(\cdot)$ pertaining to link l_2 (i.e., entries of the form $C_{l_2}[t - \tau]$ in the set $CS_{l_2}^R(\cdot)$) is maximized when $\tau_{l_2}(l_k)$ is the second-smallest of all the delay values in row k (second-smallest since, $\tau_{l_1}(l_k)$ being the smallest in row k , it can't be that $\tau_{l_2}(l_k) < \tau_{l_1}(l_k)$).

Following the structure noted above, we rearrange all the delay values in each row i in ascending order (except the delay value of 0 at position (i, i) which is left as is). We note that this structure on the delays in the table of delay values characterizes a worst-case scenario for the R policy. Finally, the structure on the table of delay values as needed in the statement of the proposition is obtained by swapping columns j and $L - j + 1$ for $j = 1, \dots, \lfloor L/2 \rfloor$, which is legitimate since the transmitters are symmetric for purposes of calculating functional evaluation complexity. \blacksquare

APPENDIX C PROOF OF PROPOSITION 3

$T_l(\cdot)$ is a function of $CS_l^R(\cdot)$. Given the structure of worst-case delay values noted in Proposition 2, we see that for

transmitter 1, the set $CS_{l_1}^R(\cdot)$ has a number of random variables that is the sum of the values in the L -tuple $(L-1, 1, 1, \dots, 1)$, since transmitter 1 has NSI of link l_1 with zero delay and hence possesses NSI of link l_1 with delays equal to all other $L-1$ values in row 1 of the delay table, and since transmitter 1 has NSI of link l_i ($2 \leq i \leq L$) with a delay that is the maximum for link l_i across all transmitters. Similarly, for transmitter 2, the set $CS_{l_2}^R(\cdot)$ has a number of random variables that is the sum of the values in the L -tuple $(1, L-1, 2, 2, \dots, 2)$, since transmitter 2 has NSI of link l_1 with a delay that is the maximum for link l_1 across all transmitters, and since transmitter 2 has NSI of link l_2 with zero delay and hence possesses NSI of link l_2 with delays equal to all other $L-1$ values in row 2 of the delay table, and since transmitter 2 has NSI of link l_i ($3 \leq i \leq L$) with a delay that is second-maximum for link l_i across all transmitters. Generalizing, we see that for transmitter i , the set $CS_{l_i}^R(\cdot)$ has a number of random variables that is the sum of the values in the L -tuple $(i-1, i-1, \dots, i-1, L-1, i, i, \dots, i)$ (where the term $L-1$ is at position i in the L -tuple). Also, given $\{C_l[t - \tau_{max}]\}_{l \in \mathcal{L}}$, the number of random variables in each of the sets $CS_{l_i}^R(\cdot)$, $1 \leq i \leq L$, reduce by 1. Therefore, the number of parameters for $T_{l_1}(\cdot)$, $T_{l_2}(\cdot)$, and $T_{l_i}(\cdot)$ are respectively $((L-1)+(L-1)-1)$, $(1+(L-1)+2(L-2)-1)$, and $((i-1)^2 + (L-1) + i(L-i) - 1)$ which simplifies to $(i(L-2) + (L-1))$. In the language of Proposition 1, these are the values of $|\mathcal{V}_1|$, $|\mathcal{V}_2|$ and $|\mathcal{V}_i|$ respectively. The required result follows immediately on appealing to Proposition 1. \square

APPENDIX D PROOF OF PROPOSITION 4

The conditional expectation in Expr. (9) is evaluated given $\mathbf{C}[t - \tau_{max}]$. Therefore, there are τ_{max} random variables in the sample path for each link l (namely, the random variables $C_l[t - (\tau_{max} - 1)]$, \dots , $C_l[t]$). Hence, the total number of random variables in the sample path is $L\tau_{max}$. Each of these random variables can take \mathcal{C} values. Consequently, the number of sample paths required in evaluating the conditional expectation in Expr. (9) is $\mathcal{C}^{L\tau_{max}}$. \blacksquare

APPENDIX E PROOF OF PROPOSITION 5

The arguments are similar to those in the proof of Proposition 1. We present it nevertheless for purposes of completeness.

Since $T_{l_i}(\cdot)$ is a function of $CS_{l_i}(\mathbf{C}[t](0 : \tau_{max}))$ and since $\{C_{l_i}[t - \tau_{l_i,max}]\}_{l_i \in \mathcal{L}}$ is made available as part of the computation of the conditional expectation (the *given*) in Eq. (11), the number of parameters to the threshold function $T_{l_i}(\cdot)$ is the number of elements in the set $\mathcal{W}_i := CS_{l_i}(\mathbf{C}[t](0 : \tau_{max})) \setminus \{C_{l_i}[t - \tau_{l_i,max}]\}_{l_i \in \mathcal{L}}$.

Each of the random variables in $CS_{l_i}(\cdot)$ can take any of the \mathcal{C} values where \mathcal{C} is the number of channel states. Thus, for $T_{l_i}(\cdot)$, there are $\mathcal{C}^{|\mathcal{W}_i|}$ different values in the domain of $T_{l_i}(\cdot)$. Each of these $\mathcal{C}^{|\mathcal{W}_i|}$ in the domain of $T_{l_i}(\cdot)$ can be independently mapped to a real number. From Proposition 13,

it suffices to consider $\mathcal{C}+1$ carefully chosen real numbers in the range of $T_{l_i}(\cdot)$. For example, from Proposition 13, for $\mathcal{C} = \{1, 2\}$, it suffices to consider the real numbers 0.5, 1.5, 2.5 as the set of possible values in the range of $T_{l_i}(\cdot)$ that each element in the domain of $T_{l_i}(\cdot)$ can be mapped to. Thus, there are $(\mathcal{C}+1)^{|\mathcal{W}_i|}$ possible choices of threshold functions for $T_{l_i}(\cdot)$. Therefore, considering all the links in the network, the total number of threshold functions that need to be considered in the domain of optimization in Expression (10) is $(\mathcal{C}+1)^{\sum_{l_i \in \mathcal{L}} \mathcal{C}^{|\mathcal{W}_i|}}$. \square

APPENDIX F PROOF OF PROPOSITION 6

By definition (see Expr. (6) and (13)), $CS^R(\mathbf{C}[t](0 : \tau_{max})) = CS(\mathbf{C}[t](0 : \tau_{max}))$. The proof is complete by considering the Expr. (7) and (12), and noting that $CS^R(\mathbf{C}[t](0 : \tau_{max}))$ and $CS(\mathbf{C}[t](0 : \tau_{max}))$ do not have CSI of link m with delay values greater than $\tau_{m,max}$, for any $m \in \mathcal{L}$, leaving the CSI of link m with delays greater than $\tau_{m,max}$ in $\mathcal{P}_{lm}^R(\mathbf{C}[t](0 : \tau_{max}))$ (and hence in $\mathcal{P}_l^R(\mathbf{C}[t](0 : \tau_{max}))$) redundant. \blacksquare

APPENDIX G PROOF OF PROPOSITION 7

We first note from Proposition 6 that, the number of random variables in the sets $CS_l^R(\cdot)$ and $CS_l(\cdot)$ are the same. The rest of the arguments are similar to those in proof of Proposition 2, and therefore omitted. \blacksquare

APPENDIX H PROOF OF PROPOSITION 8

$T_l(\cdot)$ is a function of $CS_l(\cdot)$. Following similar arguments as in the proof of Proposition 3 in Appendix C we note that, for transmitter i , the number of random variables in the set $CS_{l_i}(\cdot)$ is the sum of the values in the L -tuple $(i-1, i-1, \dots, i-1, L-1, i, i, \dots, i)$ (where the term $L-1$ is at position i in the L -tuple). Also, given $\{C_l[t - \tau_{l,max}]\}_{l \in \mathcal{L}}$, the number of random variables in each of the sets $CS_{l_i}(\cdot)$, $1 \leq i \leq L$, reduce by L . Therefore, the number of parameters for the i th threshold function $T_{l_i}(\cdot)$ is $((i-1)^2 + (L-1) + i(L-i) - L)$ which simplifies to $i(L-2)$. Thus, in the language of Proposition 5, $|\mathcal{W}_i| = i(L-2)$. Following similar arguments as in the proof of Proposition 3 in Appendix C, the total number of threshold functions that are needed to be considered in the domain of optimization in Expr. (10) is $(\mathcal{C}+1)^{\sum_{l_i \in \mathcal{L}} \mathcal{C}^{i(L-2)}}$. \blacksquare

APPENDIX I PROOF OF LEMMA 1

The proof of this lemma is similar to the proof of Lemma 4.1 in [1]. We are given that $\{A[t]\}_t$ is supportable.

This implies that there exists a policy \mathcal{F} that has the following properties: (i) \mathcal{F} uses k_i time units delayed NSI of link l_i in making its scheduling decisions at time t (where $\tau_{l,max} \leq k_i \leq t$), (ii) \mathcal{F} stabilizes the Markov chain $Z^{\mathcal{F}}[t] = \{\mathbf{Q}[t](0 : \tau_{k,max}), \mathbf{C}[t](0 : \tau_{k,max})\}$, where²⁴ $\mathbf{Q}[t](0 : \tau_{k,max}) := \{Q_{l_i}[t](0 : \tau_{k_i,max})\}_{l_i \in \mathcal{L}}$, and $\mathbf{C}[t](0 : \tau_{k,max}) := \{C_{l_i}[t](0 : \tau_{k_i,max})\}_{l_i \in \mathcal{L}}$. We note that \mathcal{F} could possibly use queue-state information in making its scheduling decisions.

First, from \mathcal{F} , we construct a new policy \mathcal{F}' , on the system state Markov chain $Y[t] = \{Q_l[t](0 : \tau_{l,max}), C[t](0 : \tau_{l,max})\}_{l \in \mathcal{L}}$ as noted next.²⁵ Consider $z \in Z^{\mathcal{F}}[t]$ where $z = \{\mathbf{q}(\tau_{l,max}+1 : k)\mathbf{q}(0 : \tau_{l,max}), \mathbf{c}(\tau_{l,max}+1 : k)\mathbf{c}(0 : \tau_{l,max})\}$, where $\mathbf{q}(\tau_{l,max}+1 : k) := \{q_{l_i}(\tau_{l_i,max}+1 : k_i)\}_{l_i \in \mathcal{L}}$, $\mathbf{q}(0 : \tau_{l,max}) := \{q_l(0 : \tau_{l,max})\}_{l \in \mathcal{L}}$, and $\mathbf{c}(\tau_{l,max}+1 : k)$ and $\mathbf{c}(0 : \tau_{l,max})$ are similarly defined. We “collapse” $Z^{\mathcal{F}}[t]$ onto $Y^{\mathcal{F}'}[t]$ as follows. Let l be the link that policy \mathcal{F} schedules in state z . Then z is collapsed onto state $y \in Y^{\mathcal{F}'}[t]$ where $y = \{\mathbf{q}(0 : \tau_{l,max}), \mathbf{c}(0 : \tau_{l,max})\}$, with \mathcal{F}' scheduling link l in state y with probability $\sum_{z'} \tilde{\pi}(z') (=: \hat{\pi}(y, l))$ where z' are all the states in $Z^{\mathcal{F}}[t]$ that “collapse” onto state y (i.e. states in $Z^{\mathcal{F}}[t]$ that match the NSI $\{\mathbf{q}(0 : \tau_{l,max}), \mathbf{c}(0 : \tau_{l,max})\}$) such that \mathcal{F} schedules link l in state z' , and where $\tilde{\pi}(z')$ is the stationary probability of being in state z' under the policy \mathcal{F} . Let $\hat{\pi}(y) = \sum_l \hat{\pi}(y, l)$. In Lemma 2, we show that the probability with which link l is scheduled in the policies \mathcal{F} and \mathcal{F}' is the same.

For $y = (\mathbf{q}, \mathbf{c})$, we let $r(y) = \Pr(\mathbf{q}|\mathbf{c})$ where $\Pr(\mathbf{q}, \mathbf{c}) = \hat{\pi}(y)$. From \mathcal{F}' , we construct a new policy \mathcal{F}_s , whose scheduling decisions are independent of all queue-state information, on the system-state Markov chain $Y[t] = \{\mathbf{Q}[t](0 : \tau_{l,max}), \mathbf{C}[t](0 : \tau_{l,max})\}$ as follows – at each time, when the delayed CSI $\mathbf{C}[t](0 : \tau_{l,max}) = \mathbf{c}$, the policy \mathcal{F}_s probabilistically schedules link l with probability $\sum_{y=(\mathbf{q},\mathbf{c})} \hat{\pi}(y, l) / \hat{\pi}(y) \times r(y)$. In Lemma 3, we show that the probability with which link l is scheduled in the policies \mathcal{F}' and \mathcal{F}_s is the same. This would then imply that \mathcal{F}_s supports $A[t]$. Since \mathcal{F}_s makes its scheduling decisions based on the delayed CSI $\mathbf{C}[t](0 : \tau_{l,max}) = \mathbf{c}$ (oblivious of queue state information), and since the service rates of all policies using the delayed CSI $\mathbf{C}[t](0 : \tau_{l,max}) = \mathbf{c}$ are considered in defining Λ , we have $E[A[t]] \in \Lambda$. ■

Lemma 2. *The policies \mathcal{F} and \mathcal{F}' schedule a link l with the same probability.*

Proof: Policy \mathcal{F} schedules link l with probability $\sum_z \tilde{\pi}(z)$, where z are all the states in $Z^{\mathcal{F}}[t]$ such that \mathcal{F} schedules link l in state z . Policy \mathcal{F}' schedules link l with probability $\sum_y \sum_{z'} \tilde{\pi}(z')$, where z' are all the states in $Z^{\mathcal{F}}[t]$ that “collapse” into state $y \in Y^{\mathcal{F}'}[t]$ such that \mathcal{F} schedules link l in state z' . This quantity is equal to $\sum_y \hat{\pi}(y, l)$ (alternatively, in state $y \in Y[t]$, the policy \mathcal{F}' schedules link l with probability $\hat{\pi}(y, l) / \hat{\pi}(y)$, and hence, overall, policy \mathcal{F}' schedules link l with probability $\sum_y \hat{\pi}(y, l) / \hat{\pi}(y) \times \hat{\pi}(y) = \sum_y \hat{\pi}(y, l)$). The

²⁴we write $\mathbf{Z}^{\mathcal{F}}[t]$ as $Z^{\mathcal{F}}[t]$ to improve legibility.

²⁵we write $\mathbf{Y}[t]$ and $\mathbf{Y}^{\mathcal{F}'}[t]$ as $Y[t]$ and $Y^{\mathcal{F}'}[t]$ respectively, to improve legibility.

proof is complete by noting that $\sum_z \tilde{\pi}(z) = \sum_y \hat{\pi}(y, l)$ since states z include all states in $Z^{\mathcal{F}}[t]$ such that \mathcal{F} schedules link l in state z , not just those that “collapse” into a particular state $y \in Y^{\mathcal{F}'}[t]$. ■

Lemma 3. *The policies \mathcal{F}' and \mathcal{F}_s schedule a link l with the same probability.*

Proof: From the proof of Lemma 2 we see that the policy \mathcal{F}' schedules link l with probability $\sum_y \hat{\pi}(y, l)$. Given that the delayed CSI $\mathbf{C}[t](0 : \tau_{l,max}) = \mathbf{c}$, the policy \mathcal{F}_s schedules link l with probability:

$$\begin{aligned} & \sum_{y=(\mathbf{q},\mathbf{c})} \frac{\hat{\pi}(y, l)}{\hat{\pi}(y)} \times r(y) \\ &= \sum_{\mathbf{q}} \frac{\hat{\pi}(y, l)}{\hat{\pi}(y)} \times r(y), \text{ where } y = (\mathbf{q}, \mathbf{c}) \\ &= \sum_{\mathbf{q}} \frac{\hat{\pi}(y, l)}{\hat{\pi}(y)} \times \Pr(\mathbf{q}|\mathbf{c}) = \sum_{\mathbf{q}} \frac{\hat{\pi}(y, l)}{\hat{\pi}(y)} \times \frac{\Pr(\mathbf{q}, \mathbf{c})}{\Pr(\mathbf{c})} \\ &= \sum_{\mathbf{q}} \frac{\hat{\pi}(y, l)}{\hat{\pi}(y)} \times \frac{\hat{\pi}(y)}{\Pr(\mathbf{c})} = \sum_{\mathbf{q}} \frac{\hat{\pi}(y, l)}{\Pr(\mathbf{c})} \end{aligned}$$

Now, averaging over all \mathbf{c} , we see that the policy \mathcal{F}_s schedules link l with probability

$$\sum_{\mathbf{c}} \sum_{\mathbf{q}} \frac{\hat{\pi}(y, l)}{\Pr(\mathbf{c})} \times \Pr(\mathbf{c}) = \sum_{(\mathbf{q},\mathbf{c})} \hat{\pi}(y, l) = \sum_y \hat{\pi}(y, l).$$
■

APPENDIX J PROOF OF THEOREM 1

Consider the optimization formulation

$$\arg \max_{\mathbf{F}(\cdot)} \sum_{l \in \mathcal{L}} Q_l[t - \tau_{l,max}] R_{l,\tau_{l,max}}(\mathbf{F}(\cdot)), \quad (17)$$

where

$$\begin{aligned} R_{l,\tau}(\mathbf{F}(\cdot)) &:= \mathbb{E}[C_l[t] F_l(\cdot)(\gamma_l + (1 - \gamma_l) \\ &\quad \prod_{m \in I_l} (1 - F_m(\cdot))) \mid \mathbf{C}[t - \tau]], \\ \mathbf{C}[t - \tau_{l,max}] &:= \{C_l[t - \tau_{l,max}]\}_{l \in \mathcal{L}}, \text{ and} \\ F_l &: \mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max})) \rightarrow \{0, 1\} \quad \forall l \in \mathcal{L}, \end{aligned}$$

with the semantics that, at time t , the transmitter node of each link l is allowed to transmit if and only if $F_l^*(\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max}))) = 1$. From Lemma 4, the optimizing solution (namely, $\mathbf{F}^*(\cdot) = \{F_l^*(\cdot)\}_{l \in \mathcal{L}}$) satisfies the threshold property $F_l^*(\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max}))) = \mathbf{1}_{\{C_l[t] \geq T_l^*(\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max})))\}}$, from Lemma 5, the optimizing solution depends only on the critical set NSI for each link l , i.e. $T_l^*(\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{l,max}))) = T_l^*(\mathcal{CS}_l(\mathbf{C}[t](0 : \tau_{l,max})))$, and from Lemma 6, the H policy supports any mean arrival rate vector λ that satisfies $(1 + \epsilon)\lambda \in \Lambda$ for any $\epsilon > 0$, but from Lemma 1, Λ is the region that encompasses all supportable arrival rates given the NSI structure in Sec. II-D, whence it follows that the H policy is throughput optimal. ■

Lemma 4. $F_l^*(\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{\max}))) = \mathbf{1}_{\{C_l[t] \geq T_l^*(\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{\max})))\}}$

Proof: The proof of this lemma is similar to the proof of Lemma 4.2 (part one) in [1], and is therefore omitted. \blacksquare

Lemma 5. $T_l^*(\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{\max}))) = T_l^*(\mathcal{CS}_l(\mathbf{C}[t](0 : \tau_{\max})))$

Proof: Even though a proof similar to the proof of Lemma 4.2 (part two) in [1] applies, we give an alternate more intuitive proof here. First, consider the NSI of link l_1 at time $t - \tau$, where $\tau < \tau_{l_1, \min}$, $\tau_{l_1, \min} = \min_{h \in \mathcal{L}} \tau_h(l_1)$. Since no transmitter has NSI of link l_1 at time $t - \tau$, none of the transmitters can make their thresholding decision based on $C_{l_1}[t - \tau]$, and therefore, $C_{l_1}[t - \tau]$ is absent from $\mathcal{CS}_{l_1}(\mathbf{C}[t](0 : \tau_{\max}))$ for all $l_1 \in \mathcal{L}$. Next, consider two links $l_i, l_j, i, j \neq 1$ such that $\tau_{l_j}(l_1) > \tau_{l_i}(l_1) + 1$ and $\tau_{l_j}(l_1)$ is the next larger delay value after $\tau_{l_i}(l_1)$ in the row corresponding to link l_1 in the table of delay values. Consider a delay value τ such that $\tau_{l_i}(l_1) < \tau < \tau_{l_j}(l_1)$, i.e. τ does not appear in the row corresponding to link l_1 in the table of delay values, or in other words, no transmitter has the NSI of link l_1 with delay τ . Now, transmitter node of link l_i does not need to make its thresholding decision based on τ since it has NSI of link l_1 with a delay smaller than τ (since $\tau_{l_i}(l_1) < \tau$) and the channel is Markovian. Therefore, $\mathcal{CS}_{l_i}(\mathbf{C}[t](0 : \tau_{\max}))$ need not contain $C_{l_1}[t - \tau]$. Lastly, consider two links $l_i, l_j, i, j \neq 1$ as above, and a delay value τ such that $\tau_{l_i}(l_1) < \tau = \tau_{l_j}(l_1)$. While it may appear that $\mathcal{CS}_{l_i}(\cdot)$ need not contain $C_{l_1}[t - \tau]$ since $\tau_{l_i}(l_1) < \tau$ and the channel is Markovian, $C_{l_1}[t - \tau]$ is nonetheless required to be in $\mathcal{CS}_{l_i}(\cdot)$ since transmitter node of link l_j makes its thresholding decision based on τ , and $C_{l_1}[t - \tau]$ helps in arbitrating when comparing the thresholding functions $T_{l_i}(\cdot)$ and $T_{l_j}(\cdot)$. From the above arguments it follows that the optimal solution does not depend on any channel state information that is not critical NSI. \blacksquare

Lemma 6. The H policy supports any mean arrival rate vector λ that satisfies $(1 + \epsilon)\lambda \in \Lambda$ for any $\epsilon > 0$.

Proof: We first define a Lyapunov function of the queue lengths available in the system state $\mathbf{Y}^{\mathcal{F}}[t]$ as follows:

$$V[t] := V[\mathbf{Y}^{\mathcal{F}}[t]] := \sum_{l \in \mathcal{L}} Q_l^2[t] \quad (18)$$

For a given arrival rate vector λ that satisfies $(1 + \epsilon)\lambda \in \Lambda$ such that $\epsilon > 0$, we show that the expected change, from one time slot to the next, in the sum of the squares of the queue lengths at all the links in the network, is negative for all but a finite number of states. This implies, from Foster's theorem (see [18]), that the system state Markov chain $\mathbf{Y}^{\mathcal{F}}[t]$ is positive recurrent, establishing that the network remains stochastically stable for this arrival rate vector.

From Equation (18), we have

$$\mathbb{E}[V[t+1] - V[t] | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]]$$

$$= \mathbb{E}[\sum_{l \in \mathcal{L}} (Q_l[t+1] - Q_l[t])(Q_l[t+1] + Q_l[t]) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]]$$

where $\mathbf{Q}[t - \tau_{\max}] = \{Q_l[t - \tau_{l, \max}]\}_{l \in \mathcal{L}}$ and $\mathbf{C}[t - \tau_{\max}] = \{C_l[t - \tau_{l, \max}]\}_{l \in \mathcal{L}}$.

Since in each time slot, the arrivals into a queue and departures out of this queue are both bounded, we have

$$\begin{aligned} & \mathbb{E}[\sum_{l \in \mathcal{L}} (Q_l[t+1] - Q_l[t])(Q_l[t+1] + Q_l[t]) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \\ & \leq K_1 + \mathbb{E}[\sum_{l \in \mathcal{L}} (Q_l[t+1] - Q_l[t])(2Q_l[t - \tau_{l, \max}]) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \quad (19) \end{aligned}$$

Next, using the queue update equation we show that

$$\begin{aligned} & \mathbb{E}[\sum_{l \in \mathcal{L}} (Q_l[t+1] - Q_l[t])(2Q_l[t - \tau_{l, \max}]) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \\ & \leq \mathbb{E}[\sum_{l \in \mathcal{L}} (R_{l, \tau_{\max}}(\mathbf{T}^*)) (2Q_l[t - \tau_{l, \max}]) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \end{aligned}$$

where \mathbf{T}^* is the optimal \mathbf{T} resulting from the optimization formulation in the scheduling policy.

Using the queue update equation (Eq. (1)), we have

$$Q_l[t+1] - Q_l[t] \leq A_l[t] + S_l[t].$$

Multiplying both sides of this inequality by $2Q_l[t - \tau_{l, \max}]$ and taking conditional expectation with respect to $\mathbf{Q}[t - \tau_{\max}]$ and $\mathbf{C}[t - \tau_{\max}]$, we have

$$\begin{aligned} & \mathbb{E}[(Q_l[t+1] - Q_l[t])(2Q_l[t - \tau_{l, \max}]) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \\ & \leq \mathbb{E}[A_{\max} 2Q_l[t - \tau_{l, \max}] | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \\ & \quad + \mathbb{E}[S_l[t] 2Q_l[t - \tau_{l, \max}] | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \quad (20) \end{aligned}$$

We note that the first term on the RHS is a constant, since the value of the term $Q_l[t - \tau_{l, \max}]$ is given in $\mathbf{Q}[t - \tau_{\max}]$. We let k_l represent this constant, and also note that for the H policy, the factor $S_l[t]$ in the second term is the same as $R_{l, \tau_{\max}}(\mathbf{T}^*)$. Therefore, we have

$$\begin{aligned} & \mathbb{E}[(Q_l[t+1] - Q_l[t])(2Q_l[t - \tau_{l, \max}]) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \\ & \leq k_l + \mathbb{E}[R_{l, \tau_{\max}}(\mathbf{T}^*) 2Q_l[t - \tau_{l, \max}] | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \quad (21) \end{aligned}$$

Summing LHS and RHS of this inequality over all l , and noting that expectation is a linear operator, we have

$$\begin{aligned} & \mathbb{E}[\sum_{l \in \mathcal{L}} (Q_l[t+1] - Q_l[t])(2Q_l[t - \tau_{l, \max}]) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \\ & \leq K_2 + \mathbb{E}[\sum_{l \in \mathcal{L}} R_{l, \tau_{\max}}(\mathbf{T}^*) 2Q_l[t - \tau_{l, \max}] | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \quad (22) \end{aligned}$$

From Inequalities (19) and (22), we have

$$\begin{aligned} & \mathbb{E}[\sum_{l \in \mathcal{L}} (Q_l[t+1] - Q_l[t])(Q_l[t+1] + Q_l[t]) | \\ & \quad \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \\ & \leq K_3 + \mathbb{E}[\sum_{l \in \mathcal{L}} R_{l,\tau_{\max}}(\mathbf{T}^*) 2Q_l[t - \tau_{l,\max}] | \\ & \quad \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \end{aligned} \quad (23)$$

That is,

$$\begin{aligned} & \mathbb{E}[V[t+1] - V[t] | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \\ & \leq K_3 + \mathbb{E}[\sum_{l \in \mathcal{L}} R_{l,\tau_{\max}}(\mathbf{T}^*) 2Q_l[t - \tau_{l,\max}] | \\ & \quad \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}]] \end{aligned} \quad (24)$$

Since $(1 + \epsilon)\lambda \in \Lambda$, there exists $\{\bar{\eta}(\mathbf{c})\}_{\mathbf{c}}$ such that

$$\bar{\eta} := \sum_{\mathbf{c} \in \mathcal{C}^L} \pi(\mathbf{c}) \bar{\eta}(\mathbf{c}) \in \Lambda, \text{ and} \quad (25)$$

$$\sum_{\mathbf{c} \in \mathcal{C}^L} \pi(\mathbf{c}) ((1 + \epsilon)\lambda_l - \bar{\eta}_l(\mathbf{c})) \leq 0 \quad \forall l \in \mathcal{L} \quad (26)$$

From the optimization formulation of the H policy, we also have

$$\begin{aligned} & \sum_{l \in \mathcal{L}} (\mathbb{E}[R_{l,\tau_{\max}}(\mathbf{T}^*) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}] = \mathbf{c}]) \\ & \quad \times (2Q_l[t - \tau_{l,\max}]) \\ & \geq \sum_{l \in \mathcal{L}} \bar{\eta}_l(\mathbf{c}) (2Q_l[t - \tau_{l,\max}]) \end{aligned} \quad (27)$$

This is because, \mathbf{T}^* is an optimal threshold, and for a given \mathbf{c} , the term $\bar{\eta}_l(\mathbf{c})$ in the RHS is the l th component of the point $\bar{\eta}(\mathbf{c}) \in \mathbb{R}^L$, and this point $\bar{\eta}(\mathbf{c})$ arises from a specific weighted combination of \mathbf{T} s (the weights being the fraction of times these \mathbf{T} s are chosen), where some of these \mathbf{T} s may not be optimal. The expectation term in the LHS results from giving a weight of 1 to \mathbf{T}^* . Also, given \mathbf{T}^* , the term $\mathbf{Q}[t - \tau_{\max}]$ in the conditional expectation on the LHS is redundant.

Let \tilde{K} be an upper bound on the LHS of Inequality (27) such that this upper bound is at least twice as large as the value of the expression on the LHS. Then, the quantity \tilde{K} minus the value of the expression on the RHS serves as an upper bound on the value of the expression on the LHS. That is,

$$\begin{aligned} & \sum_{l \in \mathcal{L}} (\mathbb{E}[R_{l,\tau_{\max}}(\mathbf{T}^*) | \mathbf{Q}[t - \tau_{\max}], \mathbf{C}[t - \tau_{\max}] = \mathbf{c}]) \\ & \quad \times (2Q_l[t - \tau_{l,\max}]) \\ & \leq \tilde{K} - \sum_{l \in \mathcal{L}} \bar{\eta}_l(\mathbf{c}) (2Q_l[t - \tau_{l,\max}]) \end{aligned} \quad (28)$$

From Inequality (26), we have

$$\sum_{\mathbf{c} \in \mathcal{C}^L} \pi(\mathbf{c}) (-\bar{\eta}_l(\mathbf{c})) \leq -(1 + \epsilon)\lambda_l \quad (29)$$

Taking expectation on both sides of Inequality (28) over $\mathbf{C}[t - \tau_{\max}]$, and noting that expectation is a linear operator, we get

$$\sum_{\mathbf{c} \in \mathcal{C}^L} \{(\mathbb{E}[\sum_{l \in \mathcal{L}} R_{l,\tau_{\max}}(\mathbf{T}^*) | \mathbf{Q}[t - \tau_{\max}], \\ \mathbf{C}[t - \tau_{\max}] = \mathbf{c}]) \times (2Q_l[t - \tau_{l,\max}])\} \pi(\mathbf{c})$$

$$\begin{aligned} & \leq \tilde{K} + \sum_{l \in \mathcal{L}} (2Q_l[t - \tau_{l,\max}]) \sum_{\mathbf{c} \in \mathcal{C}^L} \pi(\mathbf{c}) (-\bar{\eta}_l(\mathbf{c})) \\ & \leq \tilde{K} + \sum_{l \in \mathcal{L}} (2Q_l[t - \tau_{l,\max}]) (-(1 + \epsilon)\lambda_l) \quad (\text{using (29)}) \\ & \leq \tilde{K} - 2\epsilon \sum_{l \in \mathcal{L}} (Q_l[t - \tau_{l,\max}] \lambda_l) \end{aligned}$$

That is,

$$\begin{aligned} & \sum_{\mathbf{c} \in \mathcal{C}^L} \{(\mathbb{E}[\sum_{l \in \mathcal{L}} R_{l,\tau_{\max}}(\mathbf{T}^*) | \mathbf{Q}[t - \tau_{\max}], \\ & \quad \mathbf{C}[t - \tau_{\max}] = \mathbf{c}]) \times (2Q_l[t - \tau_{l,\max}])\} \pi(\mathbf{c}) \\ & \leq \tilde{K} - 2\epsilon \sum_{l \in \mathcal{L}} (Q_l[t - \tau_{l,\max}] \lambda_l) \end{aligned} \quad (30)$$

From Inequalities (24) and (30), we have

$$\begin{aligned} & \mathbb{E}[V[t+1] - V[t] | \mathbf{Q}[t - \tau_{\max}]] \\ & \leq K - 2\epsilon \sum_{l \in \mathcal{L}} Q_l[t - \tau_{l,\max}] \lambda_l \end{aligned} \quad (31)$$

Since K is a constant, the RHS in the above inequality is greater or equal to zero only for a finite number of states in the system state Markov chain $\mathbf{Y}^{\mathcal{F}}[t]$, and lesser or equal to zero for the rest of the states in $\mathbf{Y}^{\mathcal{F}}[t]$. It now follows from Foster's theorem (see [18]), that the system state Markov chain $\mathbf{Y}^{\mathcal{F}}[t]$ is positive recurrent. Hence the H policy preserves the stochastic stability of the network for the given arrival rate vector λ . Since λ is any arbitrary vector in Λ such that $(1 + \epsilon)\lambda \in \Lambda$, it follows that the H policy preserves the stochastic stability of the network for all arrival rate vectors in the interior of Λ . \blacksquare

APPENDIX K PROOF OF PROPOSITION 12

We first count the number of comparison operations in steps 4, 6, 7, 16, and 24 when the number of links in *ActiveSet* is n . We note that *ActiveSet* has L links in the first round, and that its size reduces by 1 in each subsequent round. In step 4, there are $n - 1$ comparisons for each of the n links in *ActiveSet*. Therefore, the total number of comparisons in step 4 (for all rounds together) is $L(L - 1) + (L - 1)(L - 2) + \dots + 3^2 = \mathcal{O}(L^3)$. Step 6 can be accomplished using $n - 1$ comparisons yielding a total of $(L - 1) + (L - 2) + \dots + 2 = \mathcal{O}(L^2)$ comparisons. In step 7, after suppressing the row and column of any one of the $n - 1$ links in *ActiveSet* \ { H }, we need $n - 2$ comparisons for each of the remaining $n - 1$ links in *ActiveSet* \ { K } to compute the values $\{\tau_{l,\max}\}_{l \in \text{ActiveSet} \setminus \{K\}}$, where K is the link whose row and column we suppress. We further need another $n - 1$ comparisons in the worst case to find whether eliminating this link reduces the delay for at least one channel, by comparing the $n - 1$ newly computed delays with those computed earlier in step 4. This gives us $(L - 1)((L - 1)(L - 2) + (L - 1)) + (L - 2)((L - 2)(L - 3) + (L - 2)) + \dots + 3(3.2 + 3) + 2(2.1 + 2) = \mathcal{O}(L^4)$ comparisons. In the *LC-ELDR* policy, the set *EC* has $n - 1$ elements in the worst case, and hence step 16 can be accomplished with $n - 2$ comparisons for a total of $(L - 2) + (L - 3) + \dots + 1 = \mathcal{O}(L^2)$ comparisons. In step 16

of the *LC-ERDMC* policy, the set EC has $n - 1$ elements in the worst case, and hence, as in step 7, after suppressing the row and column for each link in EC , we need $n - 2$ comparisons for each of the remaining $n - 1$ links in EC to compute the values $\{\tau_{l, \max}\}_{l \in ActiveSet \setminus \{K\}}$. We further need another $n - 1$ comparisons to find the number of links for which the newly computed delays are lesser than the corresponding delays computed earlier in step 4. This gives us $(L - 1)((L - 1)(L - 2) + (L - 1)) + (L - 2)((L - 2)(L - 3) + (L - 2)) + \dots + 3.(3.2 + 3) + 2.(2.1 + 2) = \mathcal{O}(L^4)$ comparisons. Step 24 contributes a constant to the runtime since there would only be two links left in $ActiveSet$.

Next, we count the number of multiplications and additions required in steps 5 and 24 in computing the queue-length weighted conditional expected values. Evaluation of the conditional expectation for each of the n links in $ActiveSet$ requires \mathcal{C} multiplications and \mathcal{C} additions, and one additional multiplication is required for multiplying with queue-length. Hence, in all, $L(2\mathcal{C} + 1) + (L - 1)(2\mathcal{C} + 1) + \dots + 2(2\mathcal{C} + 1) = \mathcal{O}(\mathcal{C}L^2)$ multiplications and additions are required. Therefore, the cost of comparisons, multiplications and additions together is $\mathcal{O}(\mathcal{C}L^2 + L^4)$. \blacksquare

APPENDIX L PROOF OF PROPOSITION 10

First, we need to set up some notations. Let N be the number of links in $ActiveSet$ and $T := \{1, 2, \dots, N\}$ be the set of links in $ActiveSet$ before the call to Algorithm 1 (recall that in the case of a network with complete interference, and for the first call to Algorithm 1 in the case of a network with multiple interference sets, $ActiveSet$ will have all the links in the network). Let δ_{ij} for $1 \leq i, j \leq N$ be the delay value in row i and column j of the delay table (note that $\delta_{ii} = 0 \forall i$ in our model). Let $c_{k, \tau}$ be the realization of channel-state on link l_k at time $t - \tau$. Further, let

$$\begin{aligned} \mathbf{1}\{a\} &:= \begin{cases} 1 & \text{if } a \text{ is true} \\ 0 & \text{otherwise} \end{cases} \\ \mathbf{1}\{a_i\}_{i \in I} &:= \prod_{i \in I} \mathbf{1}\{a_i\} \\ \mathbf{1}_{+}\{a_i\}_{i \in I} &:= \max \{\mathbf{1}\{a_i\}\}_{i \in I} \\ \mathbb{E}(k, \tau_k) &:= \mathbb{E}\left[C_{l_k}[t] \mid C_{l_k}[t - \tau_k] = c_{k, \tau_k}\right] \end{aligned}$$

We will call the “while” loop body from line 3 to line 23, and also the computation in lines 24–29 in Algorithm 1 in Sec. VI, a “round”. Thus, if the *LC-ELDR* policy terminates at line 26 or at line 28, then it would have executed $N - 1$ rounds (specifically, $N - 2$ rounds in the body of the “while” loop, and the last round [round $N - 1$] in lines 24–29).

Let us call the link that has not yet been eliminated from the set of contending links, an *active* link. Let $M^{(r)}$ be the active link with the largest expected data rate in round r . That is,

$$M^{(r)} := \arg \max_{k \in T \setminus \left\{ \bigcup_{1 \leq n \leq r-1} e^{(n)} \right\}} \left\{ \mathbb{E}(k, \tau_k^{(r)}) \right\},$$

where $e^{(r)}$ and $\tau_j^{(r)}$ are as defined below.

Let $e^{(r)}$ be the link other than $M^{(r)}$ that has the smallest expected data rate among all links that reduce the common delay value with which the channel state of at least one link can be accessed by the transmitter nodes of all the contending links, if this link is eliminated in round r (i.e., among all links that reduce the maximum of the delay values in a row for at least one row [corresponding to an active link] in the table of delay values [after suppressing the rows and columns corresponding to links that have been eliminated up to round r], if this link is eliminated in round r). Equivalently, $e^{(r)}$ is the link that was eliminated in round r (if r is less than the current round number), or the link that will be eliminated in round r (if r is equal to the current round number). We can write $e^{(r)}$ as in Expr. (32). The computation of $e^{(r)}$ is illustrated in Figs. 2(b) – (e) for $r = 1$, and in Figs. 2(f) – (h) for $r = 2$, for the setting in Sec. VI-A. Note that the condition “ $e^{(r)} \neq 0$ ” indicates that some link (other than the one with the largest expected data rate in round r) is eliminated in round r , and hence that the *LC-ELDR* policy can proceed to round $r + 1$ (i.e., round r is not the last round).

We let $\tau_j^{(r)}$ (not to be confused with $\tau_l(h)$ defined in Sec. II-D) be equal to $\tau_{l_j, \max}$ (see Sec. II-D) at the beginning of round r , where $\tau_{l_j, \max}$ is calculated after masking the rows and columns pertaining to the links that have been eliminated in rounds 1 to $r - 1$ (as illustrated in Figs. 2(a), 2(e), and 2(h)). That is,

$$\tau_j^{(r)} := \max_{k \in T \setminus \{j, \bigcup_{1 \leq n \leq r-1} e^{(n)}\}} \{\delta_{jk}\}$$

Note that $\tau_j^{(1)} \geq \tau_j^{(2)} \geq \dots \geq \tau_j^{(N-1)}$.

With these definitions in place, we are now ready to derive the required analytical expression for the expected saturated system throughput of the *LC-ELDR* policy. Let $i \in T$ be the link chosen by the *LC-ELDR* policy in a particular time slot, say t . We now consider the working of the *LC-ELDR* policy when it is executed at the transmitter node of link i in time slot t . From the listing of the *LC-ELDR* policy in Sec. VI, we see that link i can emerge as the “winner” (i.e., as the link chosen by the *LC-ELDR* policy for carrying transmission) in time slot t in two ways – (i) from lines 25–26 (i.e., round $N - 1$ is the last round), or (ii) from lines 8–10, 14 (i.e., an intermediate round $r < N - 1$ is the last round). Thus, link i can emerge as the “winner” when the last round is any of 1 to $N - 1$. Now, link i will emerge as the “winner” when the last round is r , $1 \leq r \leq N - 1$, if it “survives” (i.e., if it not eliminated in) round 1, survives round 2, …, survives round r .

Let us first fix a particular last round r , $1 \leq r \leq N - 1$ (i.e, the *LC-ELDR* policy runs from round 1 to round r and terminates immediately after round r). When the last round is r , if link i emerges as the “winner”, its contribution to the total expected saturated system throughput is given by

$$\sum_{\substack{c_{j, \tau_j^{(q)}} \\ j \in T, q \in \{1, 2, \dots, r\}}} \mathbb{E}\left[C_{l_i}[t] \times \right.$$

$$e^{(r)\dagger} := \begin{cases} \arg \min_{k \in T \setminus \{ \cup_{1 \leq n \leq r-1} e^{(n)}, M^{(r)} \}} \left\{ \mathbf{1}_{+} \{ \delta_{jk} > \max_{m \in T \setminus \{ j, k, \cup_{1 \leq n \leq r-1} e^{(n)} \}} \{ \delta_{jm} \} \}_{j \in T \setminus \{ k \}, \cup_{1 \leq n \leq r-1} e^{(n)}} \right. \\ \left. \times \mathbb{E}(k, \tau_k^{(r)}) \right\}, & \text{if } r < N-1, \text{ and } \mathbf{1}_{+} \{ \cdot \} \text{ above evaluates to 1 for at least one } k \in T \setminus \{ \cup_{1 \leq n \leq r-1} e^{(n)}, M^{(r)} \} \\ 0 & \text{otherwise} \end{cases} \quad (32)$$

[†]for the *LC-ELDR* policy; [§]chooses the link with the largest k if more than one link have the same minimum value

$$\begin{aligned} & \prod_{\tilde{r}=1}^{r-1} \left(\mathbf{1}\{\text{link } i \text{ is not eliminated in round } \tilde{r}\} \right. \\ & \quad \times \mathbf{1}\{\text{round } \tilde{r} \text{ is not the last round}\} \left. \right) \\ & \mathbf{1}\{\text{link } i \text{ is not eliminated in round } r\} \\ & \quad \times \mathbf{1}\{\text{round } r \text{ is the last round}\} \\ & \quad \left| \begin{array}{l} C_{l_i}[t - \tau_i^{(r)}] = c_{i, \tau_i^{(r)}}, \\ \{C_{l_i}[t - \tau_i^{(q)}] = c_{i, \tau_i^{(q)}}\}_{q \in \{1, 2, \dots, r-1\}}, \\ \{C_{l_p}[t - \tau_p^{(q)}] = c_{p, \tau_p^{(q)}}\}_{p \in T \setminus \{i\}, q \in \{1, 2, \dots, r\}} \end{array} \right] \\ & \times \Pr \{ C_{l_1}[t - \tau_1^{(1)}] = c_{1, \tau_1^{(1)}}, C_{l_1}[t - \tau_1^{(2)}] = c_{1, \tau_1^{(2)}}, \\ & \quad \dots, C_{l_N}[t - \tau_N^{(r)}] = c_{N, \tau_N^{(r)}} \} \end{aligned} \quad (33)$$

Now, consider a particular round \tilde{r} of the *LC-ELDR* policy. If link i has survived (i.e., has not been eliminated in) the first $\tilde{r}-1$ rounds, then it will survive round \tilde{r} if one of the following happens:

- 1) link i has the largest expected data rate in round \tilde{r} (hence $i = H$ and therefore $i \notin EC$; recall that we set aside the link with the largest expected data rate in each round so that it will not be eliminated in that round [see steps 6 and 7 of Algorithm 1]). We denote this condition as $s(i, \tilde{r}, 1)$.
- 2) link i does not have the largest expected data rate in round \tilde{r} (hence $i \neq H$), and link i does not reduce the maximum of the delay values in a row, for any row (corresponding to a link in *ActiveSet*) in the table of delay values (after suppressing the rows and columns corresponding to links that have been eliminated up to round \tilde{r}) if it is eliminated in round \tilde{r} (hence $i \notin EC$; see step 7 of Algorithm 1), and some link (other than the one with the largest expected data rate in round \tilde{r}) reduces the maximum of the delay values in a row, for at least one row (i.e. link) if it is eliminated in round \tilde{r} (i.e., $EC \neq \emptyset$). We denote this condition as $s(i, \tilde{r}, 2)$.
- 3) link i does not have the largest expected data rate in round \tilde{r} (hence $i \neq H$), and link i reduces the maximum of the delay values in a row, for at least one row (corresponding to a link in *ActiveSet*) in the table of delay values (after suppressing the rows and columns corresponding to links that have been eliminated up to round \tilde{r}) if it is eliminated in round \tilde{r} (hence $i \in EC$), and link i does not have the

smallest expected data rate among the links that reduce the maximum of the delay values in a row, for at least one row (i.e. link) if that link is eliminated in round \tilde{r} (hence $i \neq S$; see step 16 of Algorithm 1). We denote this condition as $s(i, \tilde{r}, 3)$.

Now, $s(i, r, 1)$ can be written mathematically as:

$$\begin{aligned} & \mathbf{1} \left\{ \mathbb{E}(i, \tau_i^{(r)}) > \max_{k \in T \setminus \{ \cup_{1 \leq n \leq r-1} e^{(n)} \}} \{ \mathbb{E}(k, \tau_k^{(r)}) \} \right\} \\ & \times \mathbf{1} \left\{ \mathbb{E}(i, \tau_i^{(r)}) \geq \max_{k \in T \setminus \{ \cup_{1 \leq n \leq r-1} e^{(n)} \}} \{ \mathbb{E}(k, \tau_k^{(r)}) \} \right\} \end{aligned} \quad (34)$$

In the expression for $s(i, r, 1)$ above, we note that splitting the comparison of the expected data rates into two – namely, (i) comparison with (the maximum of the expected data rates of) transmitters that are less than i , and (ii) comparison with transmitters that are greater than i , creates a lexicographic ordering among the transmitters. This is required to consistently resolve the “winner” in case there is a tie in the expected data rates of multiple transmitters – we always resolve in favor of the smallest numbered transmitter (as a convention) in case of a tie.

Next, $s(i, r, 2)$ can be written mathematically as:

$$\begin{aligned} & \mathbf{1}_{+} \left\{ \mathbb{E}(i, \tau_i^{(r)}) \leq \max_{k \in T \setminus \{ \cup_{1 \leq n \leq r-1} e^{(n)} \}} \{ \mathbb{E}(k, \tau_k^{(r)}) \} \right\}, \\ & \mathbb{E}(i, \tau_i^{(r)}) < \max_{k \in T \setminus \{ \cup_{1 \leq n \leq r-1} e^{(n)} \}} \{ \mathbb{E}(k, \tau_k^{(r)}) \} \\ & \times \mathbf{1} \left\{ \delta_{ji} \leq \max_{m \in T \setminus \{ i, j \}, \cup_{1 \leq n \leq r-1} e^{(n)}} \{ \delta_{jm} \} \right\}_{j \in T \setminus \{ i \}, \cup_{1 \leq n \leq r-1} e^{(n)}} \\ & \times \mathbf{1} \{ e^{(r)} \neq 0 \} \end{aligned} \quad (35)$$

Finally, $s(i, r, 3)$ can be written mathematically as:

$$\mathbf{1}_{+} \left\{ \mathbb{E}(i, \tau_i^{(r)}) \leq \max_{k \in T \setminus \{ \cup_{1 \leq n \leq r-1} e^{(n)} \}} \{ \mathbb{E}(k, \tau_k^{(r)}) \} \right\},$$

$$\begin{aligned}
\mathbb{E}(i, \tau_i^{(r)}) &< \max_{k \in T \setminus \{i\}} \max_{\substack{1 \leq n \leq r-1 \\ k > i}} \{ \mathbb{E}(k, \tau_k^{(r)}) \} \\
&\times \mathbf{1}_{\{\delta_{ji} > \max_{m \in T \setminus \{i,j\}} \max_{\substack{1 \leq n \leq r-1 \\ 1 \leq n \leq r-1}} \{ \delta_{jm} \}} \}} \max_{\substack{j \in T \setminus \{i\} \\ 1 \leq n \leq r-1}} \{ \mathbb{E}(j, \tau_j^{(n)}) \} \\
&\times \mathbf{1}_{\{e^{(r)} \neq i\}} \tag{36}
\end{aligned}$$

From (L), (34), (35) and (36), the expression for the contribution of link i to the total expected saturated system throughput when link i emerges as the “winner” when round r , $1 \leq r \leq N-1$, is the last round becomes

$$\begin{aligned}
&\sum_{\substack{c_{j, \tau_j^{(q)}} \\ j \in T \\ q \in \{1, 2, \dots, r\}}} \mathbb{E} \left[C_{l_i}[t] \prod_{\tilde{r}=1}^{r-1} \left(\mathbf{1}_{\{s(i, \tilde{r}, n)\}} \right)_{n \in \{1, 2, 3\}} \right. \\
&\times \mathbf{1}_{\{e^{(\tilde{r})} \neq 0\}} \left. \mathbf{1}_{\{s(i, r, 1)\}} \mathbf{1}_{\{e^{(r)} = 0\}} \right. \\
&\left| C_{l_i}[t - \tau_i^{(r)}] = c_{i, \tau_i^{(r)}}, \right. \\
&\left. \left\{ C_{l_i}[t - \tau_i^{(q)}] = c_{i, \tau_i^{(q)}} \right\}_{q \in \{1, 2, \dots, r-1\}}, \right. \\
&\left. \left\{ C_{l_p}[t - \tau_p^{(q)}] = c_{p, \tau_p^{(q)}} \right\}_{\substack{p \in T \setminus \{i\} \\ q \in \{1, 2, \dots, r\}}} \right] \\
&\times \prod_{m=1}^N \left(\pi_m(c_{m, \tau_m^{(1)}}) \prod_{n=1}^{r-1} p_{m, c_{m, \tau_m^{(n)}} c_{m, \tau_m^{(n+1)}}}^{(\tau_m^{(n)} - \tau_m^{(n+1)})} \right)
\end{aligned}$$

where $\pi_m(c)$ is the steady-state probability of being in state c in the channel-state Markov chain of link m , and $p_{m, ij}^{(k)}$ is the probability of reaching state j from state i in k steps in the channel-state Markov chain of link m .²⁶

Now, summing over all possible last rounds $r = 1$ to $N-1$, and over all links $i = 1$ to N , we get the expression for the expected saturated system throughput of the *LC-ELDR* policy to be

$$\begin{aligned}
&\sum_{i=1}^N \sum_{r=1}^{N-1} \sum_{\substack{c_{j, \tau_j^{(q)}} \\ j \in T \\ q \in \{1, 2, \dots, r\}}} \mathbb{E} \left[C_{l_i}[t] \prod_{\tilde{r}=1}^{r-1} \left(\mathbf{1}_{\{s(i, \tilde{r}, n)\}} \right)_{n \in \{1, 2, 3\}} \right. \\
&\times \mathbf{1}_{\{e^{(\tilde{r})} \neq 0\}} \left. \mathbf{1}_{\{s(i, r, 1)\}} \mathbf{1}_{\{e^{(r)} = 0\}} \right. \\
&\left| C_{l_i}[t - \tau_i^{(r)}] = c_{i, \tau_i^{(r)}}, \right. \\
&\left. \left\{ C_{l_i}[t - \tau_i^{(q)}] = c_{i, \tau_i^{(q)}} \right\}_{q \in \{1, 2, \dots, r-1\}}, \right. \\
&\left. \left\{ C_{l_p}[t - \tau_p^{(q)}] = c_{p, \tau_p^{(q)}} \right\}_{\substack{p \in T \setminus \{i\} \\ q \in \{1, 2, \dots, r\}}} \right] \\
&\times \prod_{m=1}^N \left(\pi_m(c_{m, \tau_m^{(1)}}) \prod_{n=1}^{r-1} p_{m, c_{m, \tau_m^{(n)}} c_{m, \tau_m^{(n+1)}}}^{(\tau_m^{(n)} - \tau_m^{(n+1)})} \right)
\end{aligned}$$

²⁶ $\pi_m(c) = \pi(c)$ and $p_{m, ij}^{(k)} = p_{ij}^{(k)}$ $\forall m \in \mathcal{L}$ in our model; see Sec. II-A

as required. \blacksquare

APPENDIX M PROOF OF PROPOSITION 11

The expression for the expected saturated system throughput of the *LC-ERDMC* policy is the same as the expression in Proposition 10, except that we need to redefine $e^{(r)}$. For that, in addition to carrying over all the notations we established in Appendix L, we need the following notations:

$$\mathbf{1}_{\infty}\{a\} := \begin{cases} 1 & \text{if } a \text{ is true} \\ \infty & \text{otherwise} \end{cases}$$

$$s_+\{a_i\}_{i \in I} := \sum_{i \in I} \mathbf{1}\{a_i\}$$

$$s_+(k, r) := s_+\{\delta_{jk} > \max_{m \in T \setminus \{j, k\}} \max_{\substack{1 \leq n \leq r-1 \\ 1 \leq n \leq r-1}} \{ \delta_{jm} \}} \}_{j \in T \setminus \{k\}} \}_{1 \leq n \leq r-1} \}_{e^{(n)}}$$

For the *LC-ERDMC* policy, we redefine $e^{(r)}$ to be the link other than $M^{(r)}$ that has the smallest expected data rate among all links that reduce the common delay value with which the channel state of a link can be accessed by the transmitters corresponding to all the contending links, for the largest number of links, if this link is eliminated in round r (i.e., among all links that reduce the maximum of the delay values in a row for the maximum number of rows [corresponding to active links] in the table of delay values [after suppressing the rows and columns corresponding to links that have been eliminated up to round r], if this link is eliminated in round r). We can write redefined $e^{(r)}$ as in Expr. (37). The required expression for the expected saturated system throughput of the *LC-ERDMC* policy is then obtained by following the rest of the arguments as in the proof of Proposition 10 in Appendix L. \blacksquare

$$e^{(r)\ddagger} := \begin{cases} \arg \min^{\ddagger} \left\{ \max_{1 \leq n \leq r-1} e^{(n)}, M^{(r)} \right\} & \left\{ \mathbf{1}_{\infty} \{ s_+(k, r) = k' \in T \setminus \left\{ \max_{1 \leq n \leq r-1} e^{(n)}, M^{(r)} \right\} s_+(k', r) \} \right. \\ & \left. \times \mathbb{E}(k, \tau_k^{(r)}) \right\}, \text{ if } r < N-1, \text{ and } \mathbf{1}_{\infty} \{ \cdot \} \text{ above evaluates to 1 for at least one } k \in T \setminus \left\{ \max_{1 \leq n \leq r-1} e^{(n)}, M^{(r)} \right\} \\ 0 & \text{otherwise} \end{cases} \quad (37)$$

[‡]for the *LC-ERDMC* policy; [§]chooses the link with the largest k if more than one link have the same minimum value
